

2019 ARGONNE LEADERSHIP  
COMPUTING FACILITY

Argonne   
NATIONAL LABORATORY

# Science

The background of the cover is a complex, abstract visualization. It features large, glowing, semi-transparent structures in shades of red, orange, and magenta. These structures are interconnected by a network of thin, blue, wireframe-like lines. Some of these lines are thicker and more prominent, while others are thinner and more numerous. The overall effect is one of a highly complex, interconnected system, possibly representing a molecular structure, a network, or a simulation of a physical process. The colors are vibrant and the lighting is dramatic, with bright highlights and deep shadows.



On the cover: A research team led by Princeton University is using the ALCF's Theta supercomputer to perform large-scale 3D simulations aimed at determining the mechanism of core-collapse supernova explosions. Depicted is the neutrino-driven roiling convection of the nuclear material behind the newly reenergized supernova shock wave that surrounds the newly birthed neutron star, just hundreds of milliseconds after shock revival. The isosurfaces are of specific entropy and the coloring follows the electron fraction. The pathlines trace infalling parcels of matter caught up in the turbulence of the explosion. The neutrino-heated turbulent bubbles help drive the supernova explosion shock wave outwards. *Image: Joseph A. Insley and Silvio Rizzi, Argonne National Laboratory; Adam Burrows and Hiroki Nagakura, Princeton University*



## 2019 SCIENCE CONTENTS

03 Science Director's Message

04 Argonne Leadership Computing Facility

06 About ALCF

07 ALCF Team

08 ALCF Computing Resources

10 Preparing for Exascale

12 Next Stop: Exascale

16 Readying Aurora for Science on Day One

## 26 Science

28 Accessing ALCF Resources for Science

30 2019 Science Highlights

31 Biological Sciences

Advancing Design and Structure Prediction of Proteins and Peptides  
[David Baker](#)

FDTD Simulations to Facilitate Early-Stage Human Cancer Detection  
[Allen Taflove](#), [Vadim Backman](#)

X-ray Microscopy of Extended 3D Objects: Scaling Towards the Future  
[Chris Jacobsen](#)

34 Computer Science

Full-State Quantum Circuit Simulation by Using Data Compression  
[Yuri Alexeev](#), [Hal Finkel](#)

35 Engineering

Analysis and Mitigation of Dynamic Stall in Energy Machines  
[Anupam Sharma](#)

Crystal Plasticity from First Principles  
[Vasily Bulatov](#)

High-Fidelity Simulations of Flow and Heat Transfer During Motored Operation of an Internal Combustion Engine  
[Paul Fischer](#)

Laminar-Turbulent Transition in Swept-Wing Boundary Layers  
[Lian Duan](#)

MaLTESE (Machine Learning Tool for Engine Simulations and Experiments)  
[Shashi Aithal](#), [Prasanna Balaprakash](#)

40 Materials Science

Advancing Electronic Stopping Simulation: From Solids to DNA  
[Yosuke Kanai](#)

Constructing and Navigating Polymorphic Landscapes of Molecular Crystals  
[Alexandre Tkatchenko](#)

Data-Driven Materials Discovery for Optoelectronic Applications  
[Jacqueline Cole](#)

Imaging and Controlling Elemental Contrast of Nanocluster in Intense X-ray  
[Phay Ho](#)

Impact of Grain Boundary Defects on Hybrid Perovskite Solar Absorbers  
[Wissam A. Saudi](#)

Optimizing Superconductor Transport Properties through Large-Scale Simulation  
[Andreas Glatz](#)

Phase Transitions in Water-Ice-Vapor System  
[Subramanian Sankaranarayanan](#)

Predictive Modeling and Machine Learning for Functional Nanoporous Materials  
[J. Ilja Siepmann](#)

48 Physics

Ab-initio Nuclear Structure and Nuclear Reactions  
[Gaute Hagen](#)

Astrophysical Particle Accelerators: Magnetic Reconnection and Turbulence  
[Dmitri Uzdensky](#)

Deep Learning at Scale for Multi-Messenger Astrophysics  
[Eliu Huerta](#)

Emulating the Universe  
[Katrin Heitmann](#)

Extreme-Scale Simulation of Supernovae and Magnetars from Realistic Progenitors  
[Sean Couch](#)

High-Fidelity Gyrokinetic Simulation of Tokamak and ITER Edge Physics  
[Choong-Seock Chang](#)

Kinetic Simulation of FRC Stability and Transport  
[Sean Dettrick](#)

PLASM-IN-SILICO: HPC Modeling of High-Intensity Laser-Solid Interaction  
[Jean-Luc Vay](#)

Scaling LHC Proton-Proton Collision Simulations in the ATLAS Detector  
[Eric Lancon](#)

Semileptonic B- and D-meson Form Factors with High Precision  
[Ruth Van de Water](#)

Simulations of Laser Experiments to Study MHD Turbulence and Non-Thermal Charged Particles  
[Petros Tzeferacos](#)

Towards a Definitive Model of Core-Collapse Supernova Explosions  
[Adam Burrows](#)

60 ALCF Projects







## SCIENCE DIRECTOR'S MESSAGE

With Mira due for retirement at the end of 2019, a final, extremely large cosmological simulation—among the five most extensive ever conducted—is set to close out its long and productive run as a leadership-class supercomputer. This work, which will be completed by harnessing the system's entirety across almost 800 million core-hours, stands to greatly benefit the cosmological community at large and promote numerous avenues of further research—including those that could advance our understanding of dark energy and dark matter. Argonne physicist Katrin Heitmann, who helped launch Mira over half a decade ago as a member of the team that performed the system's inaugural Outer Rim simulations, leads the project.

It is fitting that Mira—having been utilized to study fields ranging from biochemistry and materials science to fluid dynamics and earth science—should continue to impact science through its final simulations and beyond. I cannot help but be amazed by the fact that, right before being powered down, this system is still capable of something so useful and expansive. The research community will be taking advantage of this work for a long time.

With its 2021 arrival officially announced by the U.S. Department of Energy in mid-March, this year has seen the ALCF ramp up its efforts to deploy Aurora, which will be one of the nation's first exascale supercomputers. Conspicuous among these efforts are the upgrades being made to our data center so as to accommodate the massive system and its commensurate power and cooling requirements.

Mirroring this escalation is the productivity of our Aurora Early Science Program (ESP) projects—five simulation projects, five data projects, and five learning projects—which are instrumental to developing and shaping scientific computing in the exascale era. Each project unites computational scientists, code developers, and hardware experts to foster an environment capable of supporting emerging data and machine learning approaches in addition to traditional simulation-based scientific computing.

The Aurora ESP prepares key applications for the scale and architecture of the machine, ensuring that the research community can leverage its capabilities from the moment it enters production mode. Our staff members continue to work with experts from Intel and Cray to optimize and master its design and programming interface.

We continue to host numerous training opportunities throughout the year, including our annual Simulation, Data, and Learning Workshop, to keep our user community at the forefront of computational research. Such events help accelerate science by introducing our users to new tools and frameworks, while also helping them improve code performance and more efficiently utilize our systems.

And, ultimately, accelerating science is what the ALCF is here for.



**KATHERINE RILEY**  
ALCF Director of Science



---

## ARGONNE LEADERSHIP COMPUTING FACILITY

The ALCF enables breakthroughs in science and engineering by providing supercomputing resources and expertise to the research community.



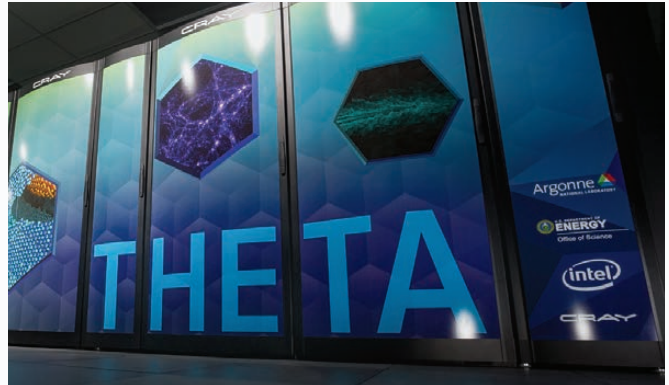




The annual ALCF Computational Performance Workshop connects facility users with staff and industry experts for hands-on assistance in boosting code performance on ALCF supercomputers.



## About ALCF



With a peak performance of more than 11 petaflops, the ALCF's Theta system is among the fastest supercomputers in the world for open scientific research.

The Argonne Leadership Computing Facility (ALCF) is a U.S. Department of Energy (DOE) Office of Science User Facility that enables breakthroughs in science and engineering by providing supercomputing resources and expertise to the research community.

ALCF computing resources—available to researchers from academia, industry, and government agencies—support large-scale computing projects aimed at solving some of the world's most complex and challenging scientific problems. Through awards of supercomputing time and support services, the ALCF enables its users to accelerate the pace of discovery and innovation across disciplines.

Supported by the DOE's Advanced Scientific Computing Research (ASCR) program, the ALCF and its partner organization, the Oak Ridge Leadership Computing Facility, operate leadership-class supercomputers that are orders of magnitude more powerful than the systems typically used for open scientific research.



# ALCF Team



The ALCF's talented and diverse staff make the facility one of the world's premier centers for scientific computing.

## Operations

The ALCF's HPC systems administrators manage and support all ALCF computing systems, ensuring users have stable, secure, and highly available resources to pursue their scientific goals. This includes the ALCF's production supercomputers, supporting system environments, storage systems, and network infrastructure. The team's software developers create tools to support the ALCF computing environment, including software for user account and project management, job failure analysis, and job scheduling.

User assistance specialists provide technical support to ALCF users and manage the workflows for user accounts and projects. In the business intelligence space, staff data architects assimilate and verify ALCF data to ensure accurate reporting of facility information.

## Science

Computational scientists with multidisciplinary domain expertise work directly with ALCF users to maximize and accelerate their research efforts. In addition, the ALCF team applies broad expertise in data science, machine learning, data visualization and analysis, and mathematics to help application teams leverage ALCF resources to pursue data-driven discoveries.

With a deep knowledge of the ALCF computing environment and experience with a wide range of numerical methods, programming models, and computational approaches, staff scientists and performance engineers help researchers optimize the performance and productivity of simulation, data, and learning applications on ALCF systems.

## Technology

The ALCF team plays a key role in designing and validating the facility's next-generation supercomputers. By collaborating with compute vendors and the performance tools community, staff members ensure the requisite programming models, tools, debuggers, and libraries are available on ALCF platforms. The team also helps manage Argonne's Joint Laboratory for System Evaluation, which houses next-generation testbeds that enable researchers to explore and prepare for emerging computing technologies.

ALCF computer scientists, performance engineers, and software engineers develop and optimize new tools and capabilities to facilitate science on the facility's current and future computing resources. This includes the deployment of scalable machine learning frameworks, in-situ visualization and analysis capabilities, data management services, workflow packages, and container technologies. In addition, the ALCF team is actively involved in programming language standardization efforts and contributes to cross-platform libraries to further enable the portability of HPC applications.

## Outreach

ALCF staff members organize and participate in training events that prepare researchers for efficient use of leadership computing systems. They also participate in a wide variety of educational activities aimed at cultivating a diverse and skilled HPC community for the future. In addition, staff outreach efforts include facilitating partnerships with industry and communicating the impactful research enabled by ALCF resources to external audiences.



# ALCF Computing Systems

## Mira

Mira is the ALCF's 10-petaflops IBM Blue Gene/Q supercomputer.

IBM Blue Gene/Q architecture	16-core, 1.6-GHz IBM PowerPC A2 processor per node	768 TB of memory
10 petaflops	49,152 nodes	5D torus interconnect
	786,432 cores	48 racks

## Cetus

Cetus is an IBM Blue Gene/Q system used to offload both debugging issues and alternative production workloads from Mira.

IBM Blue Gene/Q architecture	16-core, 1.6-GHz IBM PowerPC A2 processor per node	64 TB of memory
838 teraflops	4,096 nodes	5D torus interconnect
	65,536 cores	4 racks

## Vesta

Vesta serves at the ALCF's IBM Blue Gene/Q test and development platform.

IBM Blue Gene/Q architecture	16-core, 1.6-GHz IBM PowerPC A2 processor per node	32 TB of memory
419 teraflops	2,048 nodes	5D torus interconnect
	32,768 cores	2 racks

## Theta

Theta is the ALCF's 11.69-petaflops Intel-Cray supercomputer.

Intel-Cray XC40 architecture	4,392 nodes	70 TB of high-bandwidth memory
11.69 petaflops	281,088 cores	
64-core, 1.3-GHz Intel Xeon Phi 7230 processor per node	843 TB of memory	Aries interconnect with Dragonfly configuration
		24 racks

## Iota

Iota serves as the ALCF's Intel-Cray test and development platform.

Intel-Cray XC40 architecture	44 nodes	1 TB of high-bandwidth memory
117 teraflops	2,816 cores	
64-core, 1.3-GHz Intel Xeon Phi 7230 processor per node	12.3 TB of memory	Aries interconnect with Dragonfly configuration
		1 rack

## Cooley

Cooley is the ALCF's data analysis and visualization cluster.

Intel Haswell architecture	1 NVIDIA Tesla K80 GPU per node	47 TB of memory
293 teraflops	126 nodes	3 TB of GPU memory
Two 6-core, 2.4-GHz Intel E5-2620 processors per node	1,512 cores	FDR InfiniBand interconnect
		6 racks



# Supporting Resources

## Data Storage

At the ALCF, disk storage provides intermediate-term storage for active projects, offering a means to access, analyze, and share simulation results. Tape storage is used to archive data from completed projects.

### DISK STORAGE

The Mira system consists of 384 I/O nodes that connect to 22 storage arrays that control 13,000 disk drives with a total useable capacity of 27 PB and a maximum aggregate transfer speed of 330 GB/s over two file systems. Mira uses the General Parallel File System (GPFS) to access the storage. The Theta system consists of 30 I/O nodes that connect to a storage array that controls 2,300 disk drives with a total useable capacity of 9 PB and a maximum aggregate transfer speed of 240 GB/s. Theta uses Lustre to access this storage.

### TAPE STORAGE

The ALCF has three 10,000-slot libraries. The tape technology is currently undergoing an upgrade to replace LTO-6 tape drives with LTO-8 tape drives. The upgrade should ultimately provide up to 300 PB of effective storage (approximately five times the amount provided by the LTO-6 tapes).

## Networking

The Mira and Theta systems each have an internal proprietary network for communicating between nodes. InfiniBand enables communication between the I/O nodes and the storage system. Ethernet is used for external user access, and for maintenance and management of the systems.

The ALCF connects to other research institutions using up to 100 Gb/s of network connectivity. Scientists can transfer datasets to and from other institutions over fast research networks, such as ESnet and Internet2.

## Testbeds

Through Argonne’s Joint Laboratory for System Evaluation (JLSE), the ALCF provides access to next-generation hardware and software to explore low-level experimental computer and computational science, including operating systems, messaging, compilers, benchmarking, power measurements, I/O, and new file systems. JLSE testbeds include:

Intel Xeon Phi Knights Landing Cluster	HPE Comanche Prototype ARM64 Cluster
IBM Power System S822LC	Kubernetes Cluster with Rancher
Atos Quantum Learning Machine	NVIDIA DGX-1
Intel Xeon Platinum Skylake Cluster	IBM Elastic Storage Server GL6



---

## PREPARING FOR EXASCALE

As the future home to Aurora, one of the nation's first exascale systems, the ALCF is helping to lay the groundwork for the next generation of scientific computing.





U.S. Energy Secretary Rick Perry (left) and Intel CEO Robert Swan visited Argonne National Laboratory in March to announce the delivery of the lab's Intel-Cray exascale system, Aurora, in 2021.



# Next Stop: Exascale

The ALCF's future exascale system, Aurora, will be used to dramatically advance scientific discovery and innovation.

On March 18, 2019, the U.S Department of Energy (DOE) officially announced that Argonne National Laboratory will be home to one of the nation's first exascale systems with the delivery of Aurora in 2021.

Designed in collaboration with Intel and Cray, Aurora will help ensure continued U.S. leadership in high-end computing for scientific research.

At an announcement event at the laboratory, U.S. Secretary of Energy Rick Perry was joined by Intel CEO Robert Swan, Cray CEO Pete Ungaro, Argonne Director Paul Kearns, University of Chicago President Robert Zimmer, U.S. Senator Dick Durbin, and U.S. Representative Dan Lipinski, to underscore the significance of reaching exascale.

"Achieving exascale is imperative, not only to better the scientific community, but also to better the lives of everyday Americans," Perry said. "Aurora and the next generation of exascale supercomputers will apply HPC and AI technologies to areas such as cancer research, climate modeling, and veterans' health treatments. The innovative advancements that will be made with exascale will have an incredibly significant impact on our society."

Scientists will use Aurora to pursue some of the furthest-reaching science and engineering breakthroughs ever achieved with supercomputing. From mapping the human brain and designing new functional materials to advancing the development of alternative energy sources, Argonne's forthcoming machine will enable researchers to accelerate discoveries and innovation across scientific disciplines.

Aurora will be based on Intel's Xeon Scalable processors, high-performance Intel X<sup>e</sup> GPU compute accelerators, and Optane DC persistent memory. The system will rely on Cray's Shasta exascale-class architecture and Slingshot interconnect technology, which can provide concurrent support for advanced simulation and modeling, AI, and analytics workflows. Aurora will leverage historical advances in software investments along with increased application portability via Intel's OneAPI. The supercomputer will also introduce a new I/O system called Distributed Asynchronous Object Storage (DAOS) to meet the needs of new exascale workloads.











## AURORA SYSTEM SPECIFICATIONS

### Sustained Performance

$\geq 1$  Exaflop

### Aggregate System Memory

$> 10$  PB

### High-Performance Storage

$\geq 230$  PB,  $\geq 25$  TB/s (DAOS)

### Compute Node

Intel Xeon scalable processors,  
Xe arch-based GP-GPUs

### GPU Architecture

Xe arch-based GPU; Tile-based,  
chiplets, HBM stack, Foveros 3D  
integration

### CPU-GPU Interconnect

PCIe

### System Interconnect

Cray Slingshot; Dragonfly topology  
with adaptive routing

### Network Switch

25.6 Tb/s per switch, from 64–200 Gbs  
ports (25 GB/s per direction)

### Programming Models

Intel OneAPI, OpenMP, SYCL/DPC++

### Software Stack

Cray Shasta software stacks + Intel  
enhancements + data and learning

### Platform

Cray Shasta

*Scheduled for delivery in 2021, Aurora will feature several technological innovations to support diverse workloads that include machine learning and data-intensive tasks, in addition to traditional modeling and simulation-based research.*



# Readying Aurora for Science on Day One



ALCF researchers are working to develop computer vision and deep learning tools for processing large-scale datasets on future exascale systems.

The ALCF has continued to ramp up its efforts to prepare for Aurora's impending arrival. This includes expanding and enhancing its data center to accommodate the massive system and its power and cooling requirements, but also engaging in several activities designed to enable dramatic scientific advances in the exascale era.

The Aurora Early Science Program (ESP) is the facility's primary vehicle for ensuring the scientific computing community is ready to take advantage of the supercomputer's advanced capabilities as soon as it enters production mode. The program is designed to prepare key applications for the scale and architecture of the exascale machine, and field-test compilers and other software to pave the way for other production applications to run on the system.

Through open calls for proposals, the ESP has awarded pre-production computing time and resources to five simulation projects, five data projects, and five learning projects. The diverse set of projects reflects the ALCF's effort to create an environment that supports emerging data science and machine learning approaches alongside traditional modeling and simulation-based research.

By bringing together computational scientists, code developers, and computing hardware experts, the ESP creates a collaborative environment for optimizing applications and characterizing the behavior of the facility's future exascale system. In partnership with experts from Intel and Cray, ALCF staff members are helping train the ESP teams on the Aurora hardware design and how to program it.

This year, the Argonne-Intel Center of Excellence kicked off a series of intensive, hands-on sessions called "hackathons" to help individual ESP teams advance their efforts port and optimize applications using the Aurora software development kit.

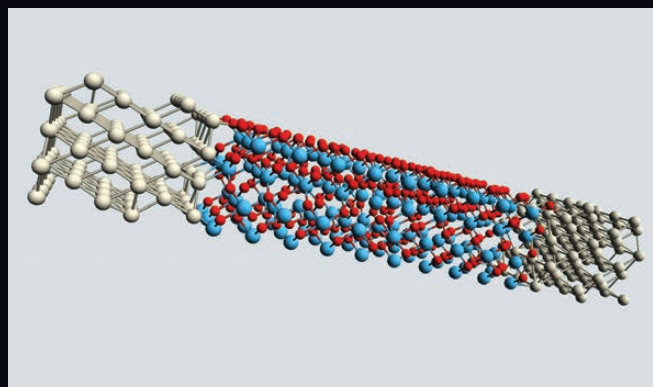
In April, the Argonne-Intel Center of Excellence held a three-day workshop for ESP teams to present a deep dive into the Aurora hardware and software environment. The event covered a wide range of topics including the Aurora software stack; its DAOS I/O system; data science/analytics workloads; programming models; performance tools; and transitioning project workflows from Theta to Aurora.

The ALCF hosted another Aurora training event in September. Open to members of the ESP and DOE's Exascale Computing Project (ECP), the three-day Aurora Programming Workshop provided attendees with additional details on the system. This event included presentations on Intel's Aurora programming models (e.g., OpenMP, SYCL/DPC++, OpenCL); open programming models (e.g., Kokkos, Raja); development hardware; and best practices and lessons learned thus far. Attendees also participated in hands-on sessions using the latest Aurora software development kit.

In addition to on-site events, the ESP continues to offer web-based tutorials to project teams on topics and tools relevant to leadership-scale computing resources, with an emphasis on data-intensive and machine learning subjects.

Together, all of these efforts are preparing the research community to harness the immense computing power of Aurora and other future exascale systems to drive a new era of scientific discoveries and technological innovations.





Hafnium oxide semiconductor with oxygen vacancies representing the oxygen leakage, inserted between platinum contacts at both ends. *Image: Olle Heinonen, Argonne National Laboratory*

## Extending Moore's Law Computing with Quantum Monte Carlo

PI Anouar Benali  
INST Argonne National Laboratory  
SFTWR QMCPACK, MKL, FFTW, BLAS/LAPACK, HDF5, ADIOS, libXML

### SCIENCE

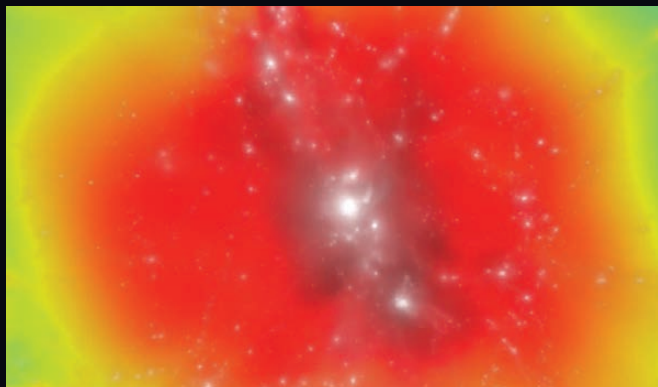
This project will carry out massive quantum Monte Carlo simulations to identify possible paths forward for extending silicon complementary metal-oxide-semiconductor (Si-CMOS)-based computing technologies beyond Moore's Law.

### DEVELOPMENT

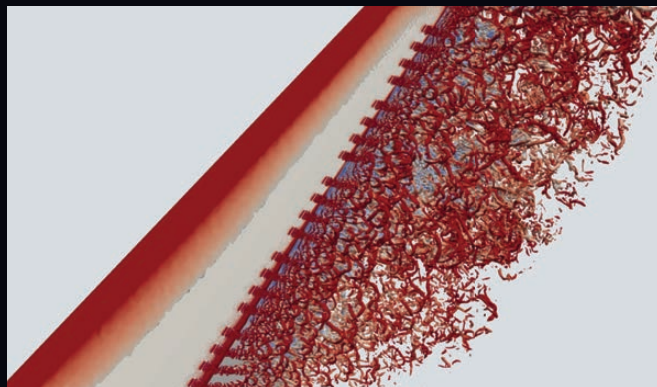
The development of QMCPACK focuses on enabling the code's kernel to run efficiently on the exascale system, while managing multi-level memory/storage hierarchy and exposing additional layers of parallelism. The team also worked on improving high-throughput internal tools to drive large campaign runs on Aurora.



## Aurora ESP | ▲ Simulation Projects



This image shows the baryon density (white) and the baryon temperature (color) of a cluster of galaxies. Image: JD Emberson and the HACC team, Argonne National Laboratory



Flow over a vertical tail/rudder assembly (grey surface) with 24 active synthetic jets which introduce vortical structures (visualized by isosurfaces of vorticity colored by local flow speed) that alter the flow, reducing separation, improving rudder performance, and thereby allowing future aircraft designs to lower drag and fuel consumption. Image: Jun Fang, Argonne National Laboratory

### Extreme-Scale Cosmological Hydrodynamics

PI      Katrin Heitmann  
INST    Argonne National Laboratory  
SFTWR   HACC, Thrust

#### SCIENCE

Researchers will use Aurora to perform cosmological hydrodynamics simulations that cover the enormous length scales characteristic of large sky surveys, while at the same time capturing the relevant small-scale physics. These simulations will help guide and interpret observations from large-scale cosmological surveys.

#### DEVELOPMENT

The team will build upon the hydrodynamics capabilities of the HACC framework to enable the modeling of observations across multiple length scales and wavebands simultaneously at high fidelity. Preparing the project's large suite of analysis tools for Aurora will be a valuable portability exercise for future system users.

### Extreme-Scale Unstructured Adaptive CFD

PI      Kenneth Jansen  
INST    University of Colorado Boulder  
SFTWR   PHASTA, PETSc, PUMI, Zoltan, parMETIS

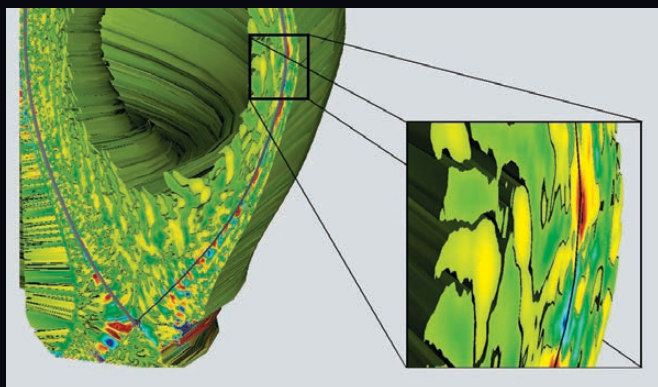
#### SCIENCE

This project will perform simulations of unprecedented scale in two complex flow regimes: aerodynamic flow control and multiphase flow. The simulations will help inform the design of next-generation aircraft and nuclear reactors by providing insights into 3D active flow control at flight scale and reactor heat exchanger flow physics, respectively.

#### DEVELOPMENT

The strong scaling needs of PHASTA will help stress test and improve Aurora's advanced interconnect. PHASTA's I/O capabilities will be enhanced to take advantage of Aurora's I/O system. This project will also make use of PETSc, PUMI, and Zoltan, preparing these libraries for future use on Aurora.





Structure of density eddies from plasma turbulence in a typical present-day tokamak edge plasma. The transition from a connected core-like turbulence eddies to isolated turbulence across the magnetic separatrix surface (black line) can be seen. Image: Dave Pugmire, Oak Ridge National Laboratory; Seung-Hoe Ku, Princeton Plasma Physics Laboratory

## High-Fidelity Simulation of Fusion Reactor Boundary Plasmas

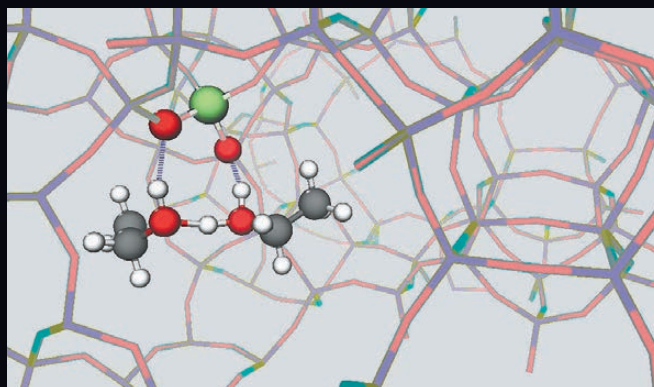
PI C.S. Chang  
 INST Princeton Plasma Physics Laboratory  
 SFTWR XGC, PETSc, PSPLINE, LAPACK, ADIOS, parMETIS

### SCIENCE

By advancing the understanding and prediction of plasma confinement at the edge, the team's simulations on Aurora will help guide fusion experiments, such as ITER, and accelerate efforts to achieve fusion energy production.

### DEVELOPMENT

XGC's problem size is expected to demand a significant fraction of the Aurora system. Computational and mathematical techniques will be developed to take advantage of Aurora's hardware structure efficiently, and to utilize memory structure for extreme-size restart data and heterogeneous physics data. This project also offers a potential comparison case for multilevel parallelization using MPI, OpenMP, and OpenACC, together with CUDA, on CPU-GPUs.



NWChemEx will provide the understanding needed to control molecular processes underlying the production of biomass. Image: Thom H. Dunning Jr., University of Washington and Pacific Northwest National Laboratory

## NWChemEx: Tackling Chemical, Materials, and Biochemical Challenges in the Exascale Era

PI Theresa Windus  
 INST Iowa State University and Ames Laboratory  
 SFTWR NWChemEx

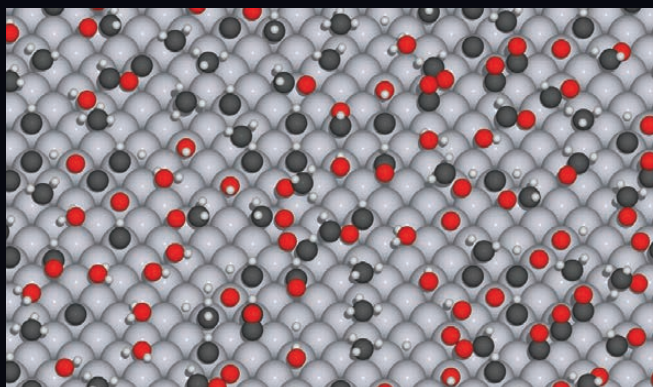
### SCIENCE

This project will use NWChemEx to address two challenges related to the production of advanced biofuels: the development of stress-resistant biomass feedstock and the development of catalytic processes to convert biomass-derived materials into fuels.

### DEVELOPMENT

The team is redesigning and reimplementing the NWChem code to enhance its scalability, performance, extensibility, and portability, positioning NWChemEx to serve as the framework for a community-wide effort to develop a comprehensive, next-generation molecular modeling package. The NWChemEx code will implement state-of-the-art algorithms for Hartree-Fock, density functional theory, and coupled cluster calculations. It will be able to effectively use multiple levels of memory as well as a partitioned global address space (PGAS) programming model to ensure that NWChemEx takes maximum advantage of the extraordinary computational capability of the Aurora exascale computing system.



AURORA ESP  
DATA PROJECTS

Data science techniques will be used in combination with quantum chemistry simulations to explore the otherwise intractable phase space resulting from gas phase molecules on catalyst surfaces to find relevant configurations and the lowest transition states between them. *Image: Eric Hermes, Sandia National Laboratories*

### Exascale Computational Catalysis

PI David Bross  
 INST Argonne National Laboratory  
 SFTWR NRRAO, RMG, PostgreSQL, Fitpy, KinBot, Sella, Balsam, NWChemEx

#### SCIENCE

Researchers will develop software tools to facilitate and significantly speed up the quantitative description of crucial gas-phase and coupled heterogeneous catalyst/gas-phase chemical systems. Such tools promise to enable revolutionary advances in predictive catalysis, crucial to addressing DOE grand challenges including both energy storage and chemical transformations.

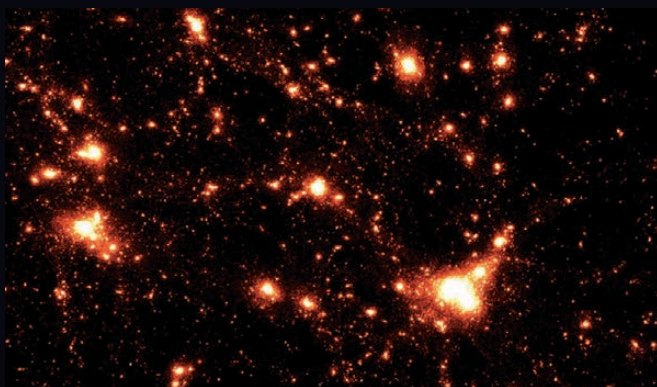
#### DEVELOPMENT

This project's diverse set of cheminformatics tools and databases will stress Aurora's deep learning software and hardware capabilities. The scientific workflow, which couples simulation, data, and learning methods, is expected to benefit from Aurora's advanced I/O hardware.

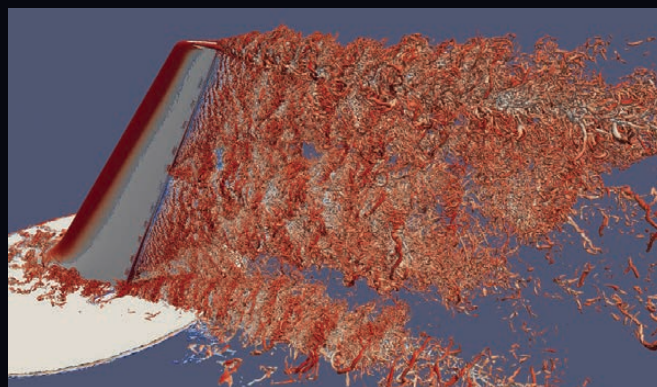
#### DATA/LEARNING METHODS

Regression, dimensionality reduction, advanced workflows, advanced statistics, reduced/surrogate models, image and signal processing, databases, and graph analytics





HACC cosmology simulations allow researchers to track the evolution of structures in great detail. Image: Joseph A. Insley, Silvio Rizzi, and the HACC team, Argonne National Laboratory



Flow visualization through an isosurface of instantaneous Q criterion colored by speed. Image: Kenneth Jansen, University of Colorado Boulder

## Dark Sky Mining

PI Salman Habib  
 INST Argonne National Laboratory  
 SFTWR Custom analyses, containers, TensorFlow, hyperparameter optimization, HACC, CosmoTools

### SCIENCE

By implementing cutting-edge data-intensive and machine learning techniques, this project will usher in a new era of cosmological inference targeted for the Large Synoptic Survey Telescope (LSST).

### DEVELOPMENT

Data mining and machine learning approaches, such as Gaussian process modeling, variational autoencoders, and Bayesian optimization, will stress the deep learning capabilities of Aurora's architecture and software. The project's end-to-end workflow is expected to benefit from the system's I/O hardware and access to databases.

### DATA/LEARNING METHODS

Classification, regression, clustering, dimensionality reduction, advanced/parallel workflows, advanced statistics, reduced/surrogate models, image and signal processing, databases, deep learning, and graph analytics

## Data Analytics and Machine Learning for Exascale Computational Fluid Dynamics

PI Kenneth Jansen  
 INST University of Colorado Boulder  
 SFTWR VTK, Paraview, PHASTA, PETSc

### SCIENCE

This project will develop data analytics and machine learning techniques to greatly enhance the value of flow simulations, culminating in the first flight-scale design optimization of active flow control on an aircraft's vertical tail.

### DEVELOPMENT

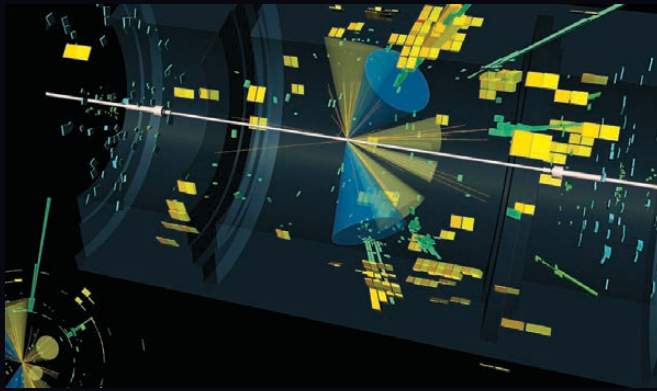
The scalable in-situ data analysis, visualization, and compression for partial differential equations involved in this work will stress Aurora's underlying deep learning software and system characteristics. The project's data-intensive workflow, which incorporates uncertainty quantification and multifidelity modeling, will benefit from the system's I/O hardware.

### DATA/LEARNING METHODS

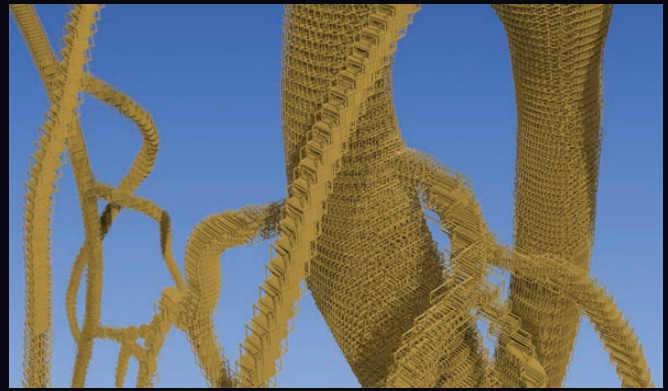
Regression, uncertainty quantification, advanced workflows, advanced statistics, reduced/surrogate models, in-situ visualization and analysis, and image and signal processing



## Aurora ESP | Data Projects



A candidate event in which a Higgs boson is produced in conjunction with top and anti-top quarks which decay to jets of particles. The challenge is to identify and reconstruct this type of event in the presence of background processes with similar signatures which are thousands of times more likely. Image: CERN



Capturing flow in the human aorta requires high-resolution fluid models. In this case, the wireframe boxes indicate each computational bounding box describing the work assigned to an individual task. Image: Liam Krauss, Lawrence Livermore National Laboratory

### Simulating and Learning in the ATLAS Detector at the Exascale

PI James Proudfoot  
 INST Argonne National Laboratory  
 SFTWR AthenaMT, Root, workflows, containers, TensorFlow

#### SCIENCE

This project will enable new physics discoveries by developing exascale workflows and algorithms that meet the growing computing, simulation and analysis needs of the ATLAS experiment at CERN's Large Hadron Collider.

#### DEVELOPMENT

The ATLAS software stack's diverse compute and analysis kernels will stress Aurora's deep learning software and hardware characteristics. An example of one key issue is the precision required by physics generators to match experimental statistics when comparing theory to experiment. The project's use of machine learning for reconstruction provides a novel use case that will likely leverage graph convolutions. This work will help to prepare Aurora for the computing needs of other large-scale experiments in high energy physics.

#### DATA/LEARNING METHODS

Classification, clustering, dimensionality reduction, advanced workflows, advanced statistics, image and signal processing, and databases

### Extreme-Scale In-Situ Visualization and Analysis of Fluid-Structure-Interaction Simulations

PI Amanda Randles  
 INST Duke University and Oak Ridge National Laboratory  
 SFTWR HARVEY, SENSEI, VTK

#### SCIENCE

The research team will develop computational models to provide detailed analysis of the role key biological parameters play in determining tumor cell trajectory in the circulatory system.

#### DEVELOPMENT

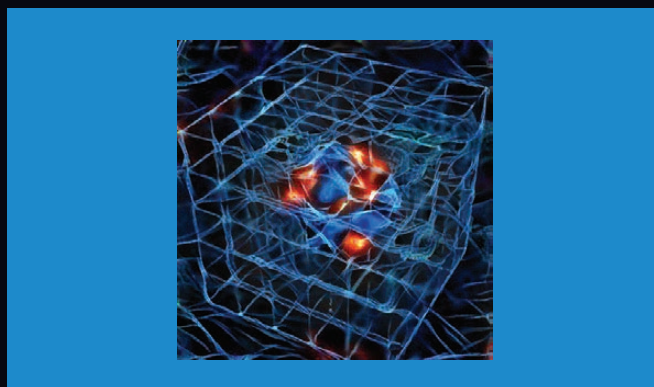
This project will stress and prepare data analysis and visualization frameworks for Aurora. Its complex, data-intensive workflows, which couple multiscale simulations and in-situ analysis, are expected to benefit from the system's I/O hardware.

#### DATA/LEARNING METHODS

Advanced workflows, reduced/surrogate models, in-situ visualization and analysis, and image and signal processing



## AURORA ESP LEARNING PROJECTS



Researchers will couple machine learning and lattice QCD simulations to advance the study of nuclei. This artistic rendering was created with [deepart.io](https://deepart.io).  
*Image: William Detmold and Phiala Shanahan, Massachusetts Institute of Technology*

### Machine Learning for Lattice Quantum Chromodynamics

---

PI William Detmold  
INST Massachusetts Institute of Technology  
SFTWR MLHMC, USQCD libraries, Spearmin, TensorFlow, HDF5, MongoDB

#### SCIENCE

By coupling advanced machine learning and state-of-the-art physics simulations, this project will provide critical input for experimental searches aiming to unravel the mysteries of dark matter while simultaneously providing insights into fundamental particle physics.

#### DEVELOPMENT

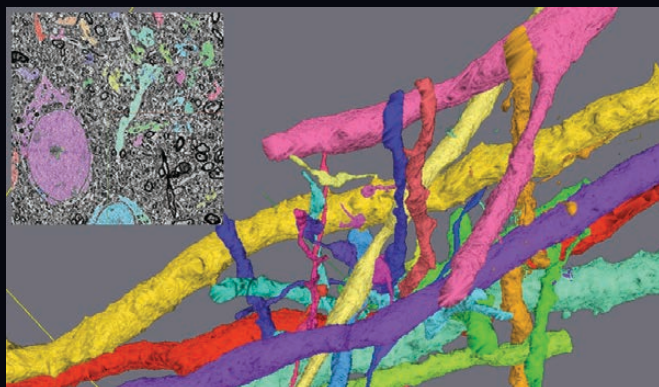
This project's novel machine learning models will stress Aurora's deep learning architecture and software stack at scale. Its workflow coupling machine learning (including deep reinforcement learning and hyperparameter optimization) and simulation could benefit from the system's I/O hardware.

#### DATA/LEARNING METHODS

Classification, regression, reinforcement learning, clustering, dimensionality reduction, advanced workflows, advanced statistics, reduced/surrogate models, and databases



## Aurora ESP | Learning Projects



Neurons rendered from the analysis of electron microscopy data. The inset shows a slice of data with colored regions indicating identified cells. Tracing these regions through multiple slices extracts the sub-volumes corresponding to anatomical structures of interest. *Image: Nicola Ferrier, Narayanan (Bobby) Kasthuri, and Rafael Vescovi, Argonne National Laboratory*

### Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience

PI Nicola Ferrier  
 INST Argonne National Laboratory  
 SFTWR TensorFlow, Horovod, flood-fill networks, AlignTK, Tomosaic

#### SCIENCE

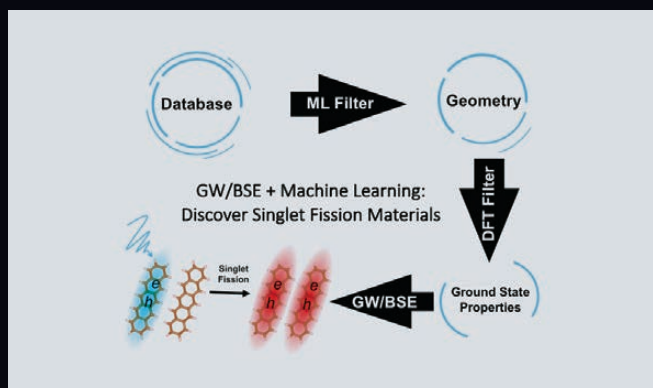
This project will develop a computational pipeline for neuroscience that will extract brain-image-derived mappings of neurons and their connections from electron microscope datasets too large for today's most powerful systems.

#### DEVELOPMENT

Compute-intensive and data-parallel deep learning models, flood-fill networks (FNN), together with large and complex data and image processing needs will stress the system's deep learning hardware/software and communication fabric. The near-real-time coupling of supercomputing resources and experiments provides a novel use case that will benefit future work on Aurora.

#### DATA/LEARNING METHODS

Classification, regression, clustering, uncertainty quantification, advanced workflows, advanced statistics, image and signal processing, and graph analytics



A schematic illustration of the project's workflow. *Image: Noa Marom, Carnegie Mellon University*

### Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials

PI Noa Marom  
 INST Carnegie Mellon University  
 SFTWR SISSO, Bayesian optimization, BerkeleyGW, Quantum Espresso

#### SCIENCE

By combining quantum-mechanical simulations with machine learning and data science, this project will harness Aurora's exascale power to revolutionize the computational discovery of new materials for more efficient organic solar cells.

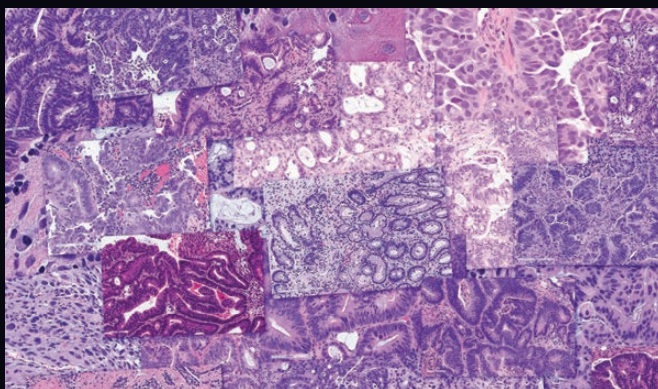
#### DEVELOPMENT

The use of feature selection and Bayesian optimization algorithms at scale will stress Aurora's deep learning stack. The project's workflow, which couples simulation, data, and learning methods, is expected to benefit from Aurora's I/O hardware.

#### DATA/LEARNING METHODS

Classification, regression, clustering, uncertainty quantification, dimensionality reduction, advanced workflows, advanced statistics, reduced/surrogate models, and databases





Predicting cancer type and drug response using histopathology images from the National Cancer Institute's Patient-Derived Models Repository. *Image: Rick Stevens, Argonne National Laboratory*

## Virtual Drug Response Prediction

PI Rick Stevens  
INST Argonne National Laboratory  
SFTWR CANDLE, Keras, TensorFlow

### SCIENCE

Utilizing large-scale data frames and a deep learning workflow, researchers will enable billions of virtual drugs to be generated and screened singly and in numerous combinations, while predicting their effects on tumor cells. This approach aims to dramatically accelerate successful drug development and provide new approaches to personalized cancer medicine.

### DEVELOPMENT

Novel models, including generative models and autoencoders, will stress the system's deep learning hardware and software stack. The project's use of data-parallel training and parallel inference, together with hyperparameter optimization, will significantly stress Aurora's communication fabric. This work also has the potential to use the system's I/O hardware to stage data for training and inference.

### DATA/LEARNING METHODS

Classification, regression, clustering, uncertainty quantification, dimensionality reduction, advanced workflows, and image and signal processing



The team's Fusion Recurrent Neural Network uses convolutional and recurrent neural network components to integrate both spatial and temporal information for predicting disruptions in tokamak plasmas. *Image: Julian Kates-Harbeck, Harvard University; Eliot Feibush, Princeton Plasma Physics Laboratory*

## Accelerated Deep Learning Discovery in Fusion Energy Science

PI William Tang  
INST Princeton Plasma Physics Laboratory  
SFTWR Fusion Recurrent Neural Net (FRNN), Keras, TensorFlow

### SCIENCE

This project will use deep learning and artificial intelligence methods to improve predictive capabilities and mitigate large-scale disruptions in burning plasmas in tokamak systems, such as ITER.

### DEVELOPMENT

The FRNN software stack will stress the system's deep learning capabilities and software stack. Large data-parallel training and hyperparameter optimization requirements will stress Aurora's communication fabric and stand to benefit from the system's I/O hardware.

### DATA/LEARNING METHODS

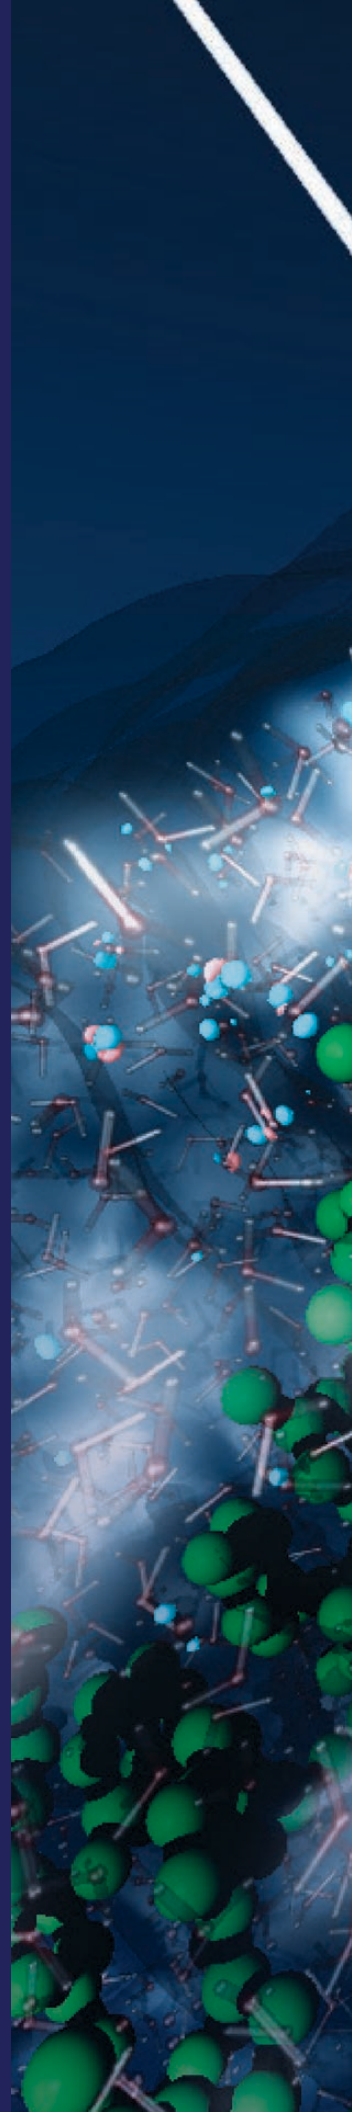
Classification, regression, clustering, and dimensionality reduction



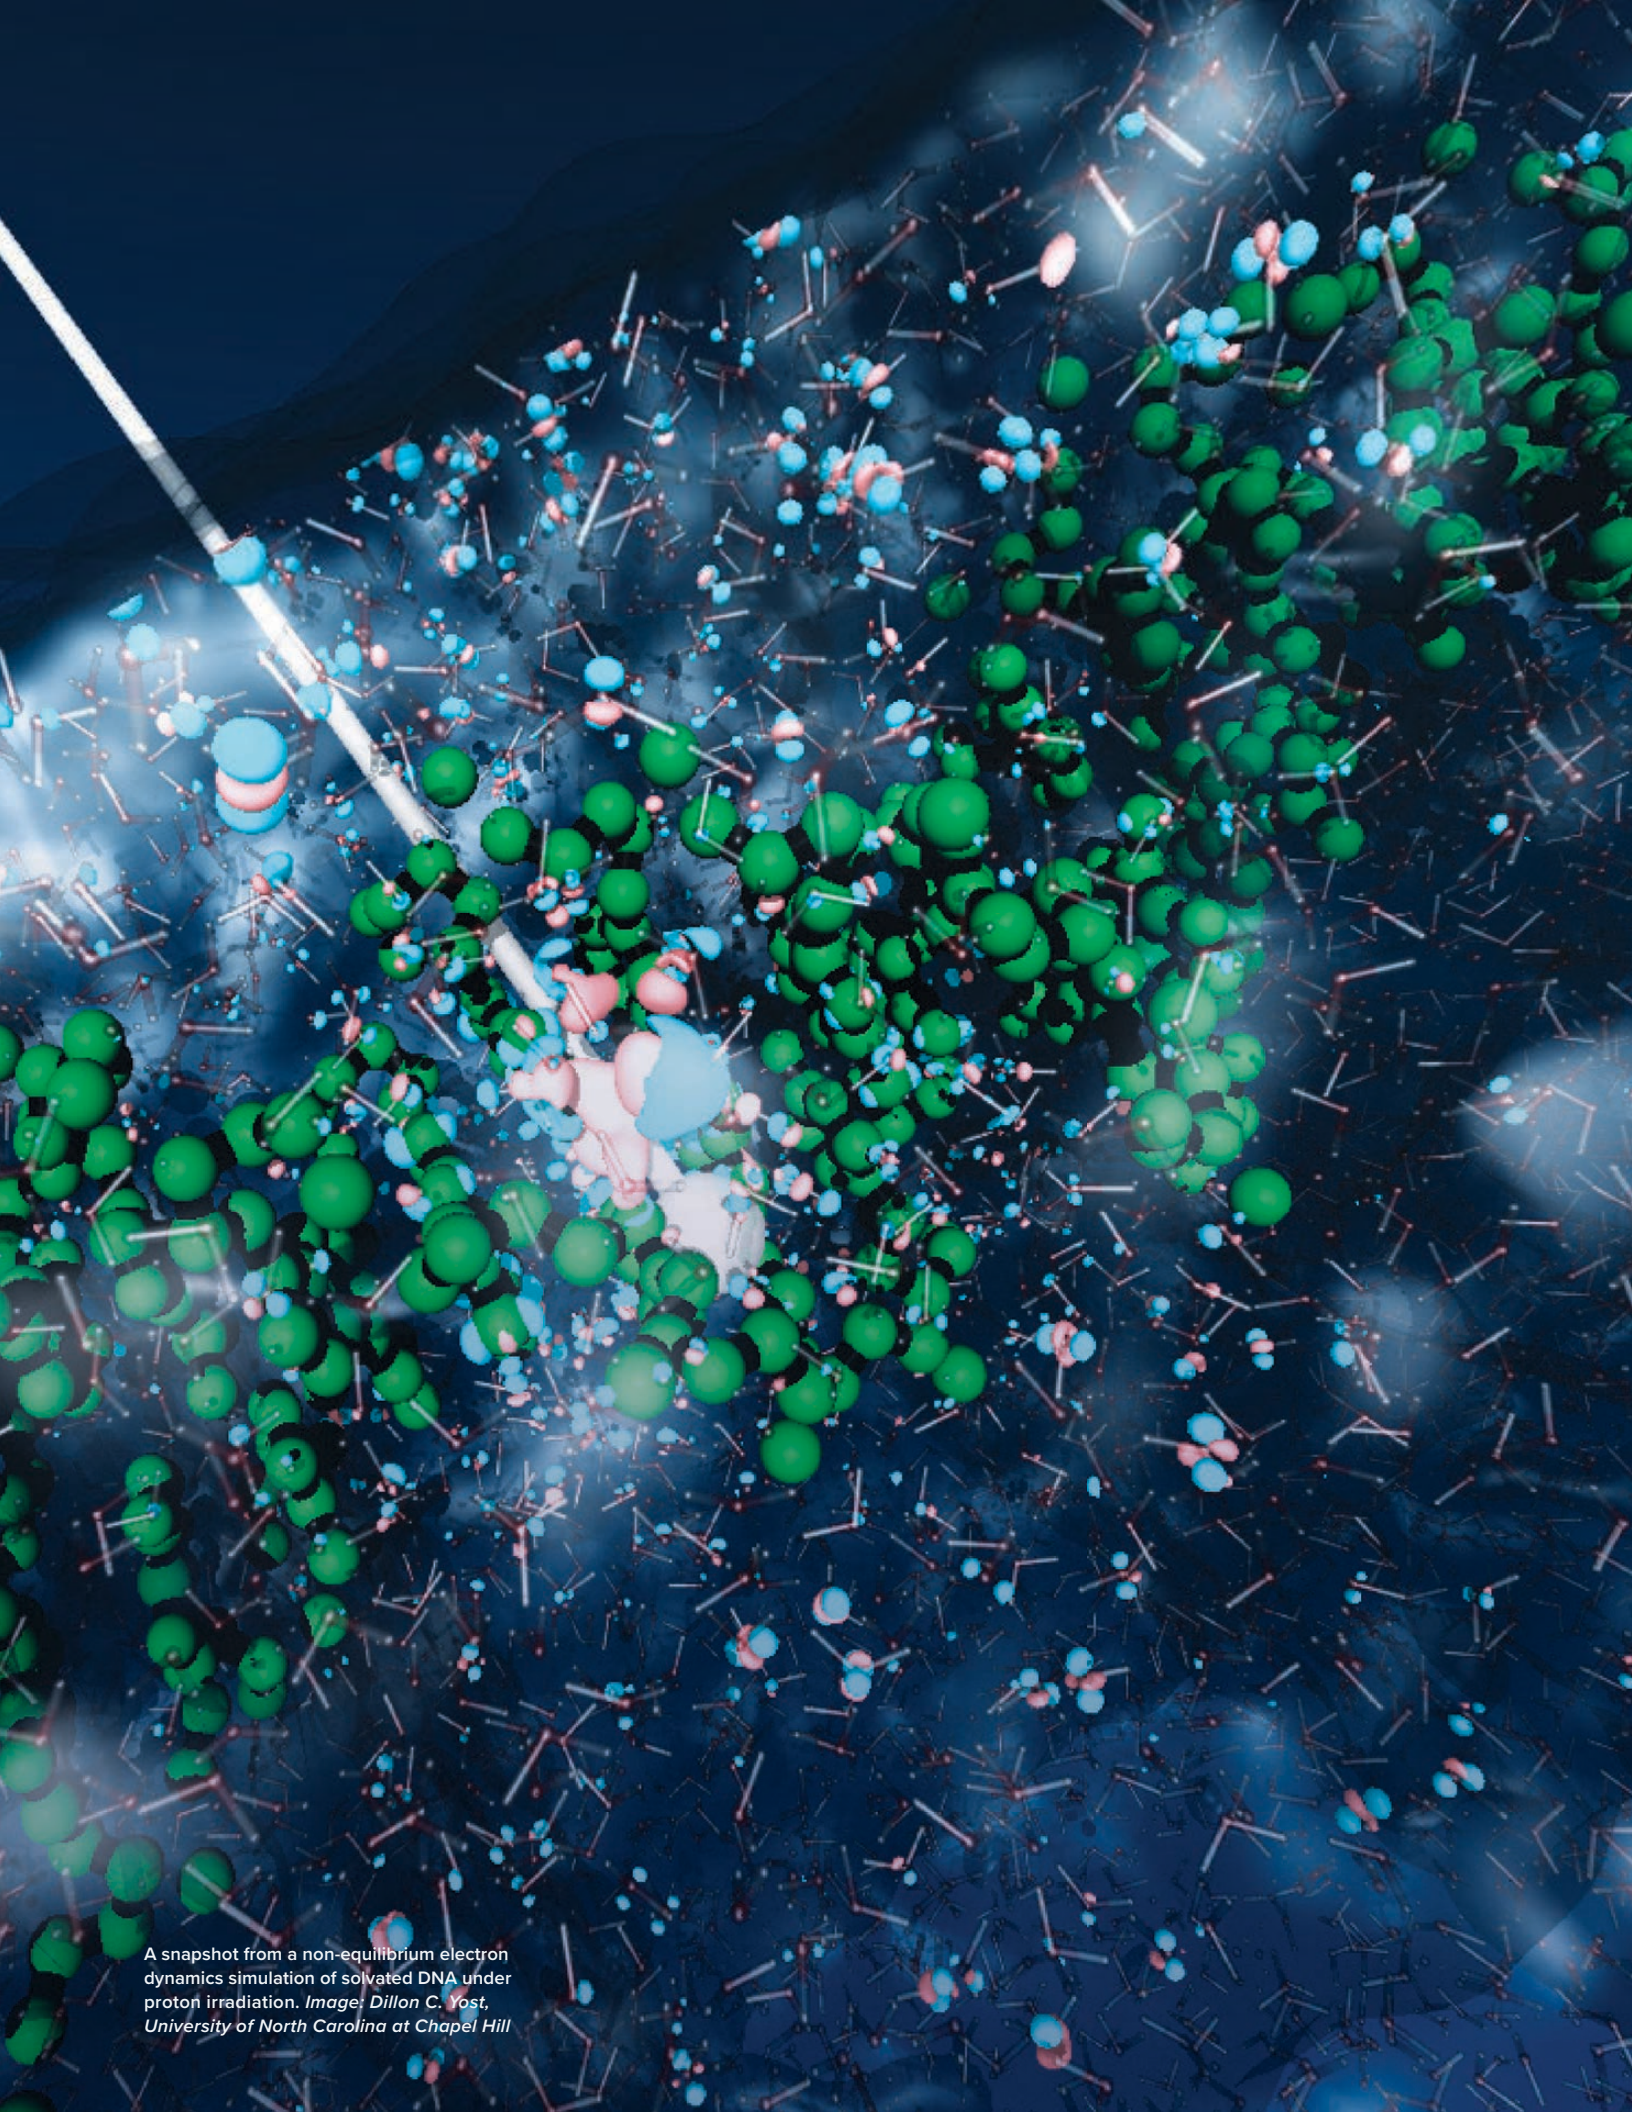
---

## SCIENCE

The ALCF is accelerating scientific discoveries in many disciplines, ranging from chemistry and engineering to physics and materials science.







A snapshot from a non-equilibrium electron dynamics simulation of solvated DNA under proton irradiation. Image: Dillon C. Yost, University of North Carolina at Chapel Hill



# Accessing ALCF Resources for Science

As a national user facility dedicated to open science, any researcher in the world with a large-scale computing problem can apply for time on ALCF computing resources.

Researchers gain access to ALCF systems for computational science and engineering projects through competitive, peer-reviewed allocations programs supported by the DOE and Argonne.

The ALCF also hosts competitive, peer-reviewed application programs designed to prepare key scientific applications and innovative computational methods for the architecture and scale of DOE supercomputers.

## Application Programs

### ADSP

The ALCF Data Science Program (ADSP) supports big data projects that require the scale and performance of leadership computing resources. ADSP projects focus on developing and improving data science techniques that will enable researchers to gain insights into very large datasets produced by experimental, simulation, or observational methods.

### ESP

As part of the process of bringing a new supercomputer into production, the ALCF conducts its Early Science Program (ESP) to prepare applications for the architecture and scale of a new system. ESP projects represent a typical system workload at the ALCF and cover key scientific areas and numerical methods.

## Allocation Programs

### INCITE

The Innovative Novel Computational Impact on Theory and Experiment (INCITE) program aims to accelerate scientific discoveries and technological innovations by awarding ALCF computing time and resources to large-scale, computationally intensive projects that address grand challenges in science and engineering.

### ALCC

The ASCR Leadership Computing Challenge (ALCC) program allocates ALCF computing resources to projects that advance the DOE mission, help to broaden the community of researchers capable of using leadership computing resources, and serve the national interests for scientific discovery, technological innovation, and economic competitiveness.

### DD

Director's Discretionary (DD) projects are dedicated to leadership computing preparation, INCITE and ALCC scaling, and efforts to maximize scientific application efficiency and productivity on leadership computing platforms.



## INCITE/ALCC BY DOMAIN

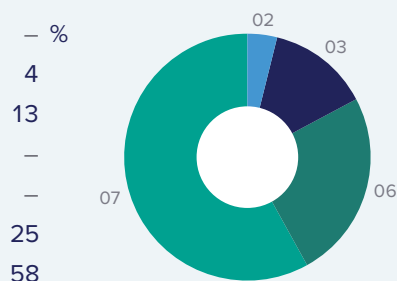
One Theta node-hour is equivalent to 13 Mira node-hours based on double-precision performance.

### 2019 INCITE

- 01 Biological Sciences
- 02 Chemistry
- 03 Earth Science
- 04 Energy Technologies
- 05 Engineering
- 06 Materials Science
- 07 Physics

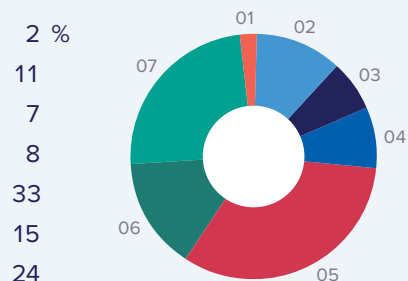
#### THETA

18M NODE HOURS



#### MIRA

220M

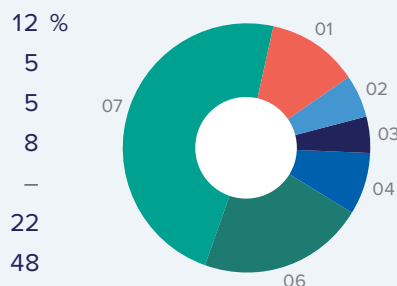


### 2019 ALCC

- 01 Chemistry
- 02 Computer Science
- 03 Earth Science
- 04 Energy Technologies
- 05 Engineering
- 06 Materials Science
- 07 Physics

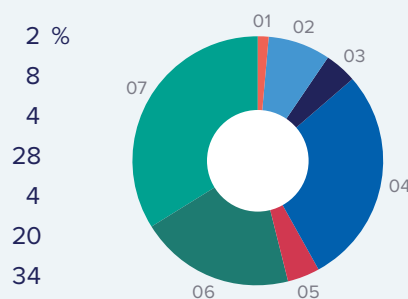
#### THETA

6M



#### MIRA

39M



ALCC data are from calendar year 2019.



# 2019 Science Highlights

The ALCF user community accelerates scientific discovery across disciplines, advancing our knowledge of the universe at all scales.

With hundreds of thousands of processors working in parallel, supercomputers allow users to achieve scientific breakthroughs that would not otherwise be possible. From detailed atomic-level simulations to massive cosmological studies, researchers can investigate extremely complex physical systems and processes that are too small or large, costly, or dangerous to study in a laboratory.

Each year, ALCF users produce impressive results, whether it be developing and demonstrating novel computational methods or publishing papers in high-impact scientific journals.

In the following pages, we present a selection of notable results from current projects supported by the ALCF's various allocation programs.

Among this year's highlights is a project that is using ALCF computing resources to advance the development of an optical microscopy technique that can predict and quantify cancer risks at extremely early stages. Another project is combining simulation and data science methods to explore novel molecular crystals with applications ranging from alternative energy materials to disease-curing pharmaceuticals.

You will also read about a collaborative effort using deep learning to classify unlabeled galaxies from the Sloan Digital Sky Survey and the Dark Energy Survey with state-of-the-art accuracy; a research team performing large-scale simulations to shed light on the complex phenomena occurring inside internal combustion engines; and several other projects pursuing breakthroughs in a wide range of studies, including supernovae, turbulent flows, peptide design, and quantum computing.



# Advancing Design and Structure Prediction of Proteins and Peptides

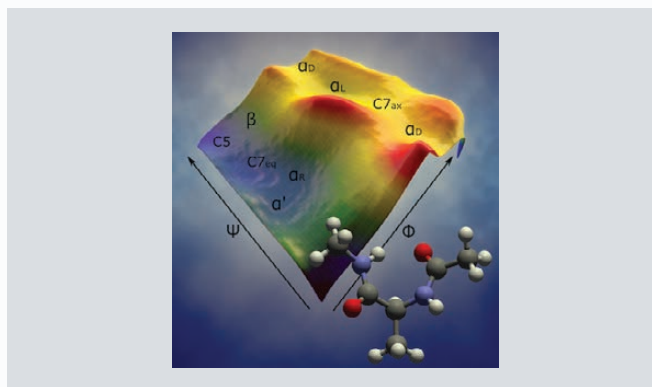
PI David Baker  
 INST University of Washington  
 AWARD INCITE  
 HOURS Mira: 5,000,000 Node-Hours;  
 Theta: 715,000 Node-Hours

**Proteins are complex molecules that drive virtually all cellular functions in living organisms. Using ALCF supercomputers, researchers from the University of Washington are advancing protein structure modeling capabilities to enable the design of synthetic mini-protein molecules, known as peptides, for the treatment of an array of diseases and medical conditions. This capability could significantly reduce the costs associated with drug design and discovery research.**

**CHALLENGE** As with the design of a protein, integral to the design of a peptide drug is creating a molecule with a rigid structure that perfectly complements the shape and charge of a binding pocket in a target site. However, for several reasons, peptide design is a more challenging task. For example, candidate molecules designed for therapeutic applications can only be rapidly explored with the development of accurate molecular force fields. This requires visualization via Ramachandran plots, difficult to obtain due to many local minima.

**APPROACH** The research team utilized GAMESS (General Atomic and Molecular Electronic Structure System) to compute the Ramachandran plots needed to visualize energetically allowed regions of the backbone dihedral angles  $\Phi$  and  $\Psi$  of amino acid residues in polypeptides. Researchers then systematically searched for energy minima in all favored conformations of alanine dipeptide by performing molecular dynamics simulations and geometry optimizations of these dipeptides at the ab-initio level of theory. ALCF staff helped implement a restrained geometry optimization protocol in GAMESS, developed scripts to generate inputs and run large numbers of calculations, and assisted with analysis and interpretation of results.

**RESULTS** The subsequent evaluation of dipeptide conformations yielded 30 unique minima, 21 of which had



The Ramachandran plot of alanine dipeptide visualizes energetically allowed regions mapped by the dihedral angles Psi and Phi (polypeptide backbone torsions) on the energy surface representing the secondary structure motifs in proteins. Image: Vikram Mulligan, Flatiron Institute

never before been reported. The evidence demonstrates the importance of some cis conformations with respect to the two peptide bonds in addition to the generally more favorable trans isomers. This was the most systematic study to locate all conformational minima of alanine dipeptide to date, with results published in the *Journal of Computational Chemistry*.

**IMPACT** The team's discovery of new energy minima provides valuable data for the development and construction of robust molecular force fields, which can be used to build more accurate models for peptide and protein design. More broadly, the design of synthetic peptides could lead to potentially revolutionary applications in medicine, manufacturing, and materials science, including the design of new therapeutics to combat disease.

## PUBLICATIONS

Mironov, V., Y. Alexeev, V. K. Mulligan, and D. G. Fedorov. "A Systematic Study of Minima in Alanine Dipeptide." *Journal of Computational Chemistry*, (December 2018), John Wiley & Sons.



# FDTD Simulations to Facilitate Early-Stage Human Cancer Detection

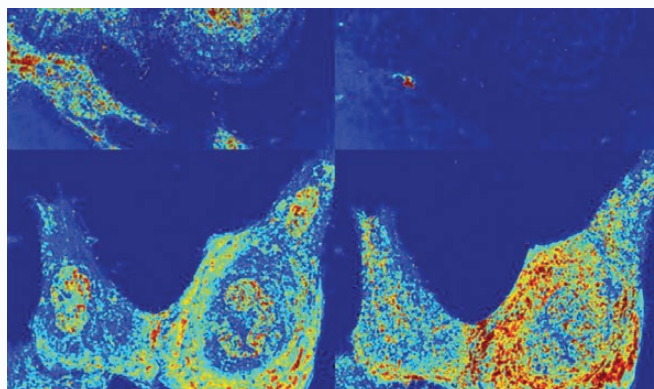
PI Allen Taflove and Vadim Backman  
 INST Northwestern University  
 AWARD INCITE  
 HOURS Mira: 5,000,000 Node-Hours

**Electrical and biomedical engineers at Northwestern University are using ALCF supercomputers to develop a low-cost, high-throughput microscopy technique capable of detecting macromolecular alterations to predict and quantify the risk of cancer at extremely early stages. The team intends to develop a technique whereby living cells can be observed to learn about their internal activity.**

**CHALLENGE** Understanding the link between intracellular motion and macromolecular structure is critical to addressing questions about macromolecular behavior in cells. The researchers have implemented finite-difference time-domain (FDTD) solution algorithms for Maxwell's equations of classical electrodynamics. These algorithms comprise the physics basis of a "microscope in a computer" that synthesizes high-resolution microscope images. This capability facilitated development of a novel optical microscopy technique, partial wave Spectroscopic (PWS) microscopy, capable of detecting the intracellular structure and macromolecular dynamics in living cells at length-scales below the diffraction limit.

**APPROACH** The researchers are employing the FDTD software package *Angora* to solve Maxwell's equations on nanoscale voxels within a simulated biological cell for an arbitrary lens-focused, apertured illumination. Integral transforms simulating refocusing lenses and apertures are applied to the FDTD-computed optical electromagnetic near field, culminating in the synthesis of full-color pixels at the image plane of the simulated microscope. These pixels are then analyzed for spectral content.

**RESULTS** As discussed in a paper published in *Nature Communications*, the researchers employed PWS to measure the intracellular nanoscale structure and macromolecular dynamics of living cells, sensitive to changes as small as 20



Alterations of the nanoscale structure of live HeLa cells after chemical fixation, as observed using PWS optical microscopy. Chemical fixation appears to alter the cellular nanoscale structure in addition to terminating its macromolecular remodeling. Image: Vadim Backman, Northwestern University

nanometers and achieving millisecond temporal resolution. FDTD computations performed in Angora were validated by comparison with experimental measurements involving nanosphere imaging phantoms. In-vitro application enabled the exploration of higher-order chromatin structure and dynamics changes attributable to cellular fixation, stem-cell differentiation, and ultraviolet (UV) irradiation. In doing so, the team discovered a new phenomenon: cellular paroxysm, a synchronous, nearly instantaneous burst of intracellular motion that occurs early during UV-induced cell death.

**IMPACT** Early-stage cancer detection has been widely recognized as one of the most critical factors to successfully treat cancer and reduce mortality. Development of a low-cost, high-throughput microscopy technique that can sense macromolecular alterations in living cells is a key step towards this goal. The high spatial-temporal resolution capability of PWS combined with a parallel FDTD implementation opens the door for high-fidelity, high-throughput early-stage cancer screening. Using the PWS technique investigated during this study, seven different human cancers are now potential candidates for early detection with minimal false positives: lung, colon, ovarian, esophageal, pancreatic, thyroid, and prostate.

## PUBLICATIONS

Gladsten, S. et al. "Multimodal Interference-Based Imaging of Nanoscale Structure and Macromolecular Motion Uncovers UV Induced Cellular Paroxysm," *Nature Communications* (April 2019), Nature Publishing Group.



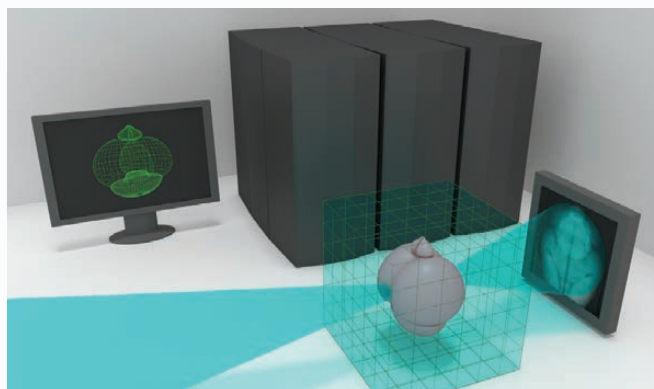
# X-ray Microscopy of Extended 3D Objects: Scaling Towards the Future

PI Chris Jacobsen  
 INST Argonne National Laboratory and Northwestern University  
 AWARD ADSP  
 HOURS Theta: 187,500 Node-Hours

The 3D x-ray microscopy of frozen hydrated biological specimens is currently approaching a limit to specimen thickness, the surpassing of which violates the pure projection approximation (PPA) needed for standard tomographic imaging. Sufficient understanding of the underlying problem has enabled the development of a novel approach to beyond-pure-projection x-ray image construction utilizing powerful—but computationally demanding—methods. This project, led by researchers from Argonne National Laboratory and Northwestern University, aims to scale up these methods to meet the challenge of high-resolution x-ray imaging beyond the PPA, benefitting not just cell and brain imaging but the full range of future nanoscale imaging activities at Argonne's Advanced Photon Source and other DOE light sources.

**CHALLENGE** Efforts to extend 3D x-ray microscopy of frozen hydrated biological specimens grow in two directions: first, sub-20-nanometer-resolution imaging of whole eukaryotic cells to elucidate their ultrastructure and its relationship with trace metal distributions, and second, 100-nanometer-resolution imaging of whole mouse brains to obtain detailed neuroanatomical maps, with an eventual goal of even higher resolution so as to resolve synaptic connections. However, both efforts have reached a limit thus far unsurpassed in x-ray microscopy: that represented by the PPA needed for standard tomographic imaging. The problem is sufficiently well understood that we can predict how an x-ray wave propagates through a complicated, extended specimen using multislice propagation, which even reproduces nanoscale x-ray waveguide effects.

**APPROACH** The research team has developed a novel approach to beyond-pure-projection x-ray image reconstruction methods whereby one constructs an approximate model of the



The approach to beyond-pure-projection x-ray image reconstruction methods involves modeling the object of study, carrying out multislice propagation, and using numerical optimization to recreate actual detected intensities at each illumination angle. Image: Ming Du and Chris Jacobsen, Argonne National Laboratory

object under study, carries out multislice propagation to a detector location, and then uses numerical optimization methods to adjust the object model to recreate the actual detected intensities at each illumination angle. The primary code for the project, running on ALCF supercomputers, is Automatic Differentiation-based Object Retrieval with dYnamic Modeling (Adorym), which uses TensorFlow, an automatic differentiation engine.

**RESULTS** The team has compared Fresnel wave propagation against finite difference methods implemented using Portable, Extensible Toolkit for Scientific Computation (PETSc), a suite of data structures and routines developed by Argonne. Automatic differentiation is being used to reconstruct objects beyond the depth of focus limit.

**IMPACT** This project develops a new computational approach to reconstructing x-ray images of thick, real-life materials. The code used will be made available to users of the Advanced Photon Source and other DOE synchrotron light sources.



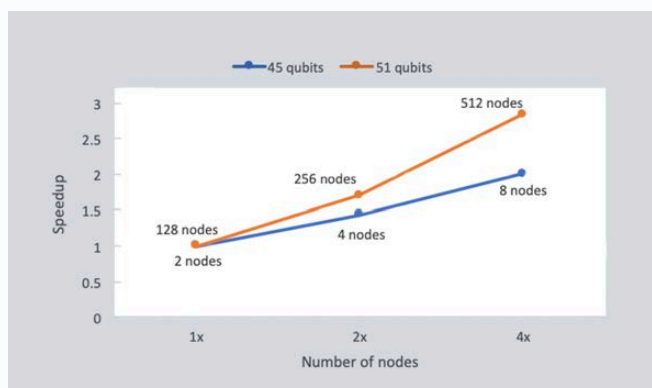
# Full-State Quantum Circuit Simulation by Using Data Compression

PI Yuri Alexeev, Hal Finkel  
 INST Argonne National Laboratory  
 AWARD DD  
 HOURS Theta: 250,000 Node-Hours

Quantum circuit simulations are critical tools for evaluating and validating the design of new quantum algorithms and machines. The size of such simulations, however, is limited by the memory capacity of current HPC systems. A research team from Argonne National Laboratory and the University of Chicago has developed a new quantum circuit simulation technique that leverages data compression, trading computation time and fidelity to reduce the memory requirement of full-state quantum circuit simulations.

**CHALLENGE** The path toward building noisy intermediate-scale quantum (NISQ) machines, such as upcoming IBM's 58-qubit quantum computer and Google's 72-qubit quantum computer, will require intermediate-scale quantum circuit simulators to calibrate and verify the hardware. Unfortunately, today's practical full-simulation limit is 47 qubits. The reason is that the number of quantum state amplitudes required for a full simulation increases exponentially with the number of qubits, making available physical memory the limiting factor to store all amplitudes.

**APPROACH** To simulate general circuits with higher qubit count and depth, the team developed a hybrid solution that combines their tailored lossy compression method with adaptive error bounds at each timestep of the simulation. The hybrid approach optimizes for compression speed and makes sure that errors due to lossy compression are uncorrelated (an important property for comparing simulation output with physical machines). This new compression technique is implemented on the full-state quantum circuit simulator Intel-QS with Argonne-developed lossy compression package SZ. The general simulation framework was then run on the ALCF's Theta supercomputer. To enable larger simulations in the future, researchers are working to port



Scalability of Intel-QS/SZ code on Theta for 45 and 51 qubit simulations to apply the Hadamard gate to each gate. Image: Yuri Alexeev, Argonne National Laboratory.

and optimize Intel-QS and other simulators on the ALCF's forthcoming exascale system, Aurora.

**RESULTS** The team's approach reduces the memory requirement of quantum circuit simulations, allowing for larger quantum system simulations with the same memory capacity. Experiments show that the approach reduces the memory requirement of simulating the 61-qubit Grover's search algorithm from 32 exabytes to just 768 terabytes of memory (compression factor of about 42,000 times) on Theta using aggregated memory of 4,096 nodes. These results suggest that the techniques can increase the simulation size by two to 16 qubits for general quantum circuits while maintaining 0.985 output fidelity on average. A paper detailing these findings has been accepted at the SC19 technical program.

**IMPACT** Classical simulation of quantum circuits is crucial for better understanding the operations and behaviors of quantum computation. By reducing the memory footprint of such simulations, the team's novel approach provides researchers and developers with a platform for quantum software debugging and hardware validation for modern quantum devices that have more than 50 qubits. The developed simulator will be used by Argonne and partnering institutions for a variety of DOE projects in areas such as basic energy sciences, high energy physics, fusion energy sciences, and advanced scientific computing.

## PUBLICATIONS

Wu, X.-C., S. Di, F. Cappello, H. Finkel, Y. Alexeev, and F. T. Chong. "Full-State Quantum Circuit Simulation by Using Data Compression," *SC19: International Conference for High Performance Computing, Networking, Storage and Analysis* (November 2019).



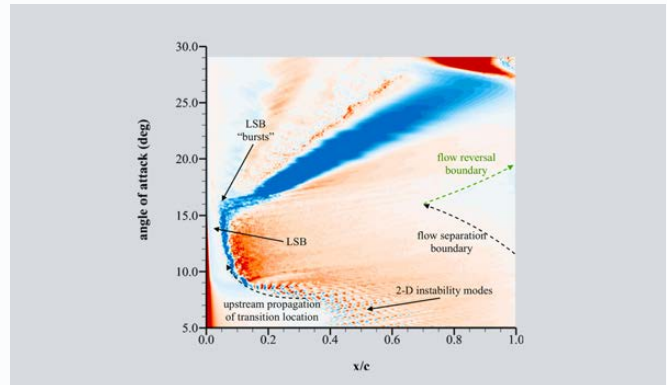
# Analysis and Mitigation of Dynamic Stall in Energy Machines

PI Anupam Sharma  
 INST Iowa State University  
 AWARD ALCC  
 HOURS Mira: 3,218,750 Node-Hours

**Dynamic stall in rotating blades of power generation turbomachines such as wind turbines results in large unsteady loads that can lead to fatigue and/or catastrophic failure. In addition, dynamic stall causes increased acoustic emissions (noise) and reduced aerodynamic performance of these machines. While the general characteristics of dynamic stall have long been known, the mechanisms of stall onset are not well understood. These mechanisms involve unsteady boundary layer flow physics, which requires high-fidelity numerical computations for resolving the relevant spatial and temporal scales. Iowa State University researchers are using wall-resolved large eddy simulations to investigate boundary layer flow physics near stall incipience, which holds the key to developing efficient stall-mitigation techniques.**

**CHALLENGE** Dynamic stall is a common occurrence in wind turbines that operate in the highly turbulent atmospheric boundary layer with large spatial and temporal gradients. The complex, unsteady boundary-layer behavior at stall onset holds the key to understanding and subsequently mitigating dynamic stall, but it has not been investigated due to the extreme computational demands required to achieve the adequate spatial and temporal resolution. This makes the use of leadership-class computing systems a necessity for any rigorous, in-depth study.

**APPROACH** The researchers are using wall-resolved large-eddy simulations to study the fluid dynamic processes that occur in the boundary layer near stall incipience and to investigate novel ideas for the passive mitigation of dynamic stall in wind and gas turbine applications. Several mechanisms, including wavy leading edge, vortex generators, pressure-side tabs (Gurney flaps), and back-flow flaps are being evaluated for their ability to mitigate dynamic stall. The backbone of the



Boundary layer flow features analyzed using a space-time plot of skin-friction coefficient on the suction surface of a pitching airfoil. The onset of dynamic stall, which occurs when the airfoil angle of attack exceeds 16 degrees, is marked by the bursting of the laminar separation bubble (LSB). Image: Anupam Sharma, Iowa State University; Miguel Visbal, Air Force Research Laboratory

researchers' computational framework is the high-order accurate flow solver FDL3DI, which solves the full Navier-Stokes equations. It has been demonstrated to scale linearly on up to 32,000 cores of the ALCF's Mira supercomputer.

**RESULTS** The researchers have investigated the effect of airfoil thickness on the onset of dynamic stall, as detailed in a paper published in the *Journal of Fluid Mechanics*. They are continuing their investigation of the impact of sinusoidal leading-edge modifications to airfoil, which has been found to lessen the intensity of stall.

**IMPACT** This project's systematic numerical investigation will enhance our understanding of what triggers dynamic stall. This, in turn, will help advance the development of robust, efficient, and quiet energy-conversion machines. Passive mitigation of stall, also explored in this work, is a desirable alternative to active given its comparative simplicity, economy, and robustness.

## PUBLICATIONS

Sharma, A., and M. Visbal. "Numerical Investigation of the Effect of Airfoil Thickness on Onset of Dynamic Stall," *Journal of Fluid Mechanics* (July 2019), Cambridge University Press.

Selvaraj, S., and A. Sharma. "Effect of Leading Edge Serrations on Dynamic Stall at  $Re=30,000$ ," *AIAA SciTech Forum* (January 2019), San Diego, California, American Institute of Aeronautics and Astronautics.



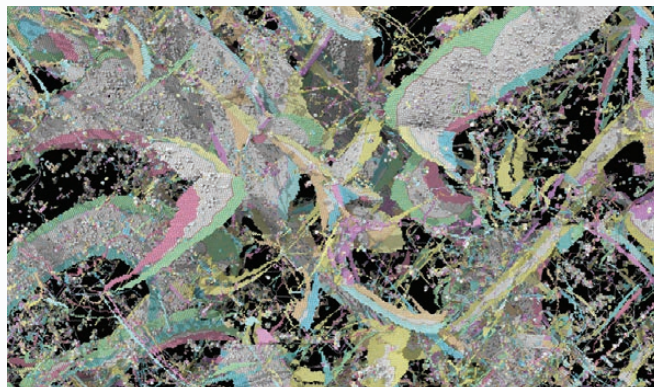
# Crystal Plasticity from First Principles

PI Vasily Bulatov  
INST Lawrence Livermore National Laboratory  
AWARD INCITE  
HOURS Mira: 17,000,000 Node-Hours

This project, led by researchers from Lawrence Livermore National Laboratory, uses large-scale molecular dynamics (MD) simulations to gain understanding of material strength and other technologically relevant mechanical properties. This research aims to settle two longstanding controversies in classical physical metallurgy: the microscopic origin of staged strain hardening and the nature of dislocation patterns.

**CHALLENGE** Prediction of crystal plasticity and strength directly from the underlying dynamics of atomic motion has been beyond the limits of existing computing abilities. It is critically important to be able to look inside the material while it is being strained to settle which of the existing theories are correct or to have sufficient data to develop and support new models if needed. At present, in-situ microscopy observations are possible only in thin electron-transparent films, where neither strain hardening nor dislocation patterns are observed. MD simulations are currently the only means to permit, in principle, simultaneous mechanical testing of bulk crystal plasticity in silico and fully detailed in-situ observations of the underlying atomic dynamics. Because of their immense computational cost—necessitating use of leadership-class computing systems—direct MD simulations of crystal plasticity had been regarded as impossible. However, by applying a newly established simulation capability, the team has demonstrated that direct cross-scale MD simulations (i.e., simulations large enough to represent macroscopic crystal plasticity and yet sufficiently detailed to trace every bit of atomic motion) of plasticity and strength of aluminum and tantalum metals are feasible.

**APPROACH** Leveraging ALCF supercomputers, the team's approach combines massive MD simulations and detailed on-the-fly and post-processing analyses and characterizations



Large-scale MD simulations of strained aluminum metal performed on Mira produce dense networks of dislocation. Image: Alexander Stukowski, Technische Universität Darmstadt

of underlying events in the life of atoms and dislocations that, taken together, define crystal plasticity response.

**RESULTS** The researchers have developed network analysis tools through the creation of a special library of Dislocation Network Analysis (DNA) functions that can be combined to perform arbitrary topological analyses. These DNA functions are being tested on networks of increasing size.

In addition, they successfully ran a molecular dynamics simulation with 24 billion atoms and are now extracting dislocation networks to search for dislocation patterns. Their hope is to soon simulate dislocation pattern formation in a model with 100 billion tantalum atoms.

**IMPACT** Ultimately, the computations performed will provide definitive data on the origin of staged strain hardening and the nature of dislocation patterns, while also increasing our understanding of material strength and other technologically relevant mechanical properties. Furthermore, the use of cross-scale simulations foreshadows the development of high-performance computing in the coming years.



# High-Fidelity Simulations of Flow and Heat Transfer During Motored Operation of an Internal Combustion Engine

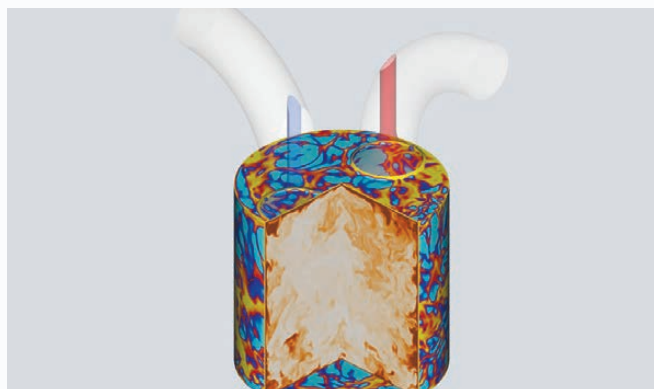
PI Paul Fischer  
 INST Argonne National Laboratory  
 AWARD ALCC  
 HOURS Theta: 468,750 Node-Hours

The gas motion in the interior of an internal combustion engine has a strong impact on the engine's efficiency and pollutant emissions through its effect on mixing, combustion, and heat transfer. To inform the design of future engines, researchers from the Swiss Federal Institute of Technology, Argonne National Laboratory, and the Aristotle University of Thessaloniki, Greece, are performing large-scale simulations on ALCF computing resources to shed light on the complex phenomena occurring inside internal combustion engines.

**CHALLENGE** Current understanding of stochastic in-cylinder processes is inadequate or limited and thus prevents engine designers from achieving significantly higher efficiency. However, with high-performance computing resources and accurate, science-based computational fluid dynamics (CFD) simulations, researchers can improve the fundamental understanding of complex combustion phenomena to establish and characterize the physical causes of stochastic events.

**APPROACH** The research team developed a workflow to automate the use of multiple grids and exploit the accurate grid-to-grid interpolation of Argonne's massively parallel CFD code, Nek5000. Running on the ALCF's Theta system, they simulated the gas-exchange process over multiple cycles of the Transparent Combustion Chamber (TCC-III) engine (for which extensive validation data are available from the University of Michigan). The results obtained with these highly accurate simulations will help clarify ambiguities in the interpretation of experimental measurements and will complement commercial engine simulation codes by providing improved benchmark-type simulations.

**RESULTS** Building upon previous work, the team has completed a simulation campaign to study the evolution of turbulence inside the optically accessible, single-cylinder



Instantaneous temperature distribution inside the cylinder and heat flux along the cylinder walls during the compression stroke in the TCC-III engine. Image: George Giannakopoulos, ETH Zurich

TCC-III engine. This effort involved carrying out high-order, wall-resolved large-eddy simulations using the spectral element method to investigate the gas-exchange process and in-cylinder flow evolution. The simulations allowed the team to quantify the in-cylinder tumble generation and breakdown, and investigate its impact on the evolution of momentum and thermal boundary layers inside the cylinder. The capabilities of modern HPC architectures combined with the solver's high granularity allowed for a significant reduction in the wall-time needed to simulate one cycle, while simultaneously increasing the resolution by at least an order of magnitude compared to the current state-of-the-art in scale resolving engine simulations.

**IMPACT** By improving our understanding of the turbulent flow, heat transfer, and combustion processes taking place inside internal combustion engines, this work will help advance the development of more efficient and environmentally friendly vehicles. In addition, the team aims to redefine the capability limits of high-fidelity simulations, paving the way for detailed investigations of more complex phenomena as increasingly powerful supercomputers become available.

## PUBLICATIONS

Giannakopoulos, G. K., C. E. Frouzakis, P. F. Fischer, A. G. Tomboulides, and K. Boulouchos. "LES of the Gas-Exchange Process Inside an Internal Combustion Engine Using a High-Order Method," *Flow, Turbulence and Combustion* (accepted for publication), Springer Science and Business Media LLC.

Giannakopoulos, G. K., C. E. Frouzakis, K. Boulouchos, P. F. Fischer, and A. G. Tomboulides. "Direct Numerical Simulation of the Flow in the Intake Pipe of an Internal Combustion Engine," *International Journal of Heat and Fluid Flow* (December 2017), Elsevier.



Engineering |  Simulation

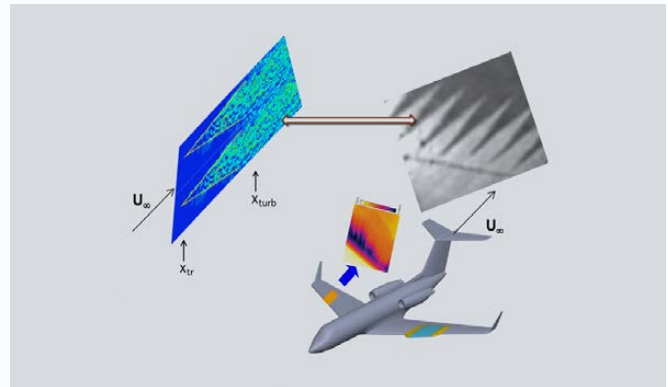
# Laminar-Turbulent Transition in Swept-Wing Boundary Layers

PI Lian Duan  
 INST Missouri University of Science and Technology  
 AWARD INCITE  
 HOURS Mira: 6,000,000 Node-Hours

Skin-friction drag accounts for approximately one half of the total drag for long-haul transport aircraft. Due to the substantial increase in drag in the course of laminar-turbulent transition, delay of crossflow-influenced transition in swept-wing boundary layers via laminar flow technology is a leading contender for reducing the aircraft fuel burn. Researchers from The Ohio State University and NASA Langley Research Center are harnessing ALCF supercomputers to enable simulations of the entire crossflow transition process—from the laminar regime to the fully turbulent regime—over a transonic natural-laminar-flow wing with high chord Reynolds numbers relevant to the transport aircraft.

**CHALLENGE** In particular, the work will focus on generating a detailed knowledge base that will address the existing gaps in crossflow transition prediction, including (1) the excitation and control of crossflow instabilities via discrete roughness elements, and (2) the laminar breakdown of crossflow vortices preceding the onset of turbulent flow and the properties of developed turbulence in the post-translational region.

**APPROACH** Laminar-turbulent boundary layer transition, including the receptivity and breakdown phases, is studied via direct numerical simulations (DNS) based on the notoriously intractable Navier-Stokes equations. Full resolution of the receptivity and breakdown phases necessitates use of leadership-class computing systems. The primary code employed is the high-order finite-difference solver HyperWENO. HyperWENO solves the compressible Navier-Stokes equations in generalized curvilinear coordinates that describe the evolution of the flow's density, momentum, and total energy. ALCF staff members have been instrumental in efforts to scale HyperWENO for use on Theta and Mira.



DNS captures swept-wing boundary layer transition over a transonic aircraft with laminar flow technology. Near-wall flow visualizations confirm the sawtooth nature of transition front as a generic feature of transition due to stationary crossflow vortices regardless of the type of secondary instability. *Image: NASA*

**RESULTS** The team has completed multiple DNS to investigate the receptivity stage of boundary-layer transition and is preparing larger-scale simulations that include stages of laminar breakdown and fully turbulent flows.

**IMPACT** This project will significantly advance our ability to predict and control laminar-turbulent transitions in 3D boundary layers. It will also help delineate the limitations of the current predictive tools, providing a basis to develop alternate models for cases in which those tools are either inapplicable or lead to unacceptable errors. Beyond aerospace engineering, the work will more broadly impact the design of devices involving 3D transitional and turbulent boundary layers, such as wind turbines and turbomachinery.



# MaLTESE (Machine Learning Tool for Engine Simulations and Experiments)

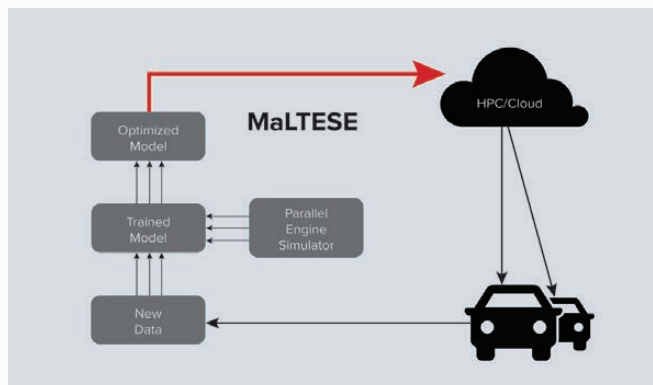
PI Shashi Aithal, Prasanna Balaprakash  
 INST Argonne National Laboratory  
 AWARD DD  
 HOURS Theta: 2,000 Node-Hours

Despite advances in vehicle technologies, continuously optimal engine operation during the course of a driver's commute has not yet been possible. Researchers at Argonne National Laboratory are developing a deep learning framework called MaLTESE (Machine Learning Tool for Engine Simulations and Experiments) that can predict engine performance and emissions for an entire commute in real time, thereby allowing for the potential on-board optimization and control in autonomous and connected vehicles.

**CHALLENGE** There are over 30 different design and engine operating parameters that affect the fuel economy, emissions, and engine performance of a vehicle. This large number of parameters, combined with the variability of individual driving habits, makes continuous engine optimization via onboard diagnostic tools currently infeasible.

**APPROACH** MaLTESE couples a fast, robust, physics-based, real-time engine simulator called pMODES (Parallel Multi-fuel Otto Diesel Engine Simulator) with a deep neural network (DNN). The research team used MaLTESE to concurrently simulate the performance and emissions of the half-hour commutes of 250,000 Chicago rush hour drivers in real time. This utilized nearly the full capacity of Theta and generated massive amounts of drive-cycle related information such as efficiency, fuel consumption, greenhouse gases, and emissions.

Statistical and machine-learning techniques were then employed to distill the information from this massive dataset to deconvolute the impact of the input variables on engine performance and emissions. A subset of the drive-cycle data was then used to train the DNN to "learn" the relationship between various input parameters and vehicle performance (torque/power) along with emissions, allowing for future



Conceptual use of MaLTESE and high-performance computing for autonomous vehicles. Image: Shashi Aithal, Argonne National Laboratory

prediction of drive-cycle characteristics of various driving scenarios.

**RESULTS** The team's simulation on Theta is considered the single largest drive-cycle simulation conducted concurrently in real time and also the first machine learning-based prediction of drive-cycle characteristics of thousands of cars on city roads and freeways during rush hour. The trained DNN model was able to predict drive-cycle characteristics to within 1 percent of those computed by the simulator but about four orders of magnitude faster, putting machine learning-based engine operation prediction within the realm of on-board real-time adaptive learning and control. These findings were presented at the ISC High Performance 2019 conference.

**IMPACT** The ultimate goal of MaLTESE is to enable better design and optimization of connected vehicles for increased fuel efficiency, reduced emission, better engine performance, and longer engine life. In addition to helping automotive manufacturers overcome technical hurdles in development and maturation of autonomous/connected vehicles, MaLTESE could allow automotive manufacturers and regulatory agencies to examine various "what-if" scenarios in understanding fuel economy and emissions at an unprecedented scale.

## PUBLICATIONS

Aithal, S. M., P. Balaprakash. "MaLTESE: Large-Scale Simulation-Driven Machine Learning for Transient Driving Cycles." *ISC High Performance 2019. Lecture Notes in Computer Science*, vol 11501. Springer, Cham.



# Advancing Electronic Stopping Simulation: From Solids to DNA

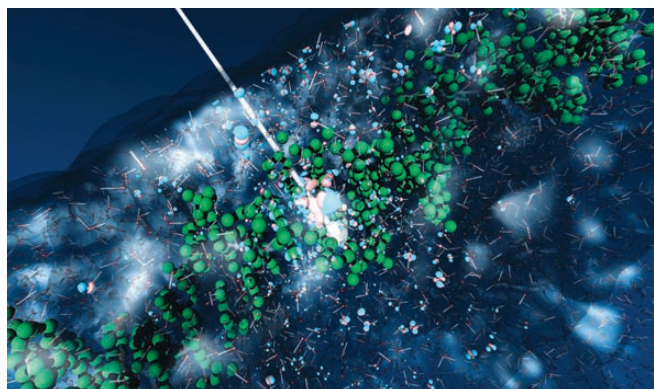
PI Yosuke Kanai  
 INST University of North Carolina at Chapel Hill  
 AWARD INCITE  
 HOURS Mira: 4,190,000 Node-Hours,  
 Theta: 1,586,000 Node-Hours

Electronic stopping refers to the dynamical transfer of kinetic energy from energetic charged particles (e.g. protons) to electrons in a target matter, consequently inducing massive electronic excitations therein. Elucidation of this phenomenon as it occurs in various systems under ion irradiation contributes to impactful breakthroughs in a number of modern technologies. A team of researchers are using predictive simulations to model electronic stopping dynamics in DNA due to their importance in various applications such as proton-beam cancer therapy.

**CHALLENGE** Predictive modeling of electronic stopping processes has remained a great challenge for many decades because of the difficulties involved in accurately describing the quantum-mechanical excitation of electrons; however, recent innovations in first-principles calculation methodologies and massively parallel supercomputers have enabled accurate simulations of these processes at the atomistic level. The researchers now seek to further advance simulation capabilities so as to model complex systems like solvated DNA under various ion irradiations.

**APPROACH** This project continues to develop a highly-scalable implementation of real-time, time-dependent density functional theory using the Qbox/Qb@ll code. Hundreds of thousands of processors in the Mira and Theta systems are used to simulate the quantum-mechanical electronic response of complex systems. For instance, the simulation of solvated DNA in water involves over 13,000 electrons, representing a enormous computational challenge.

**RESULTS** As detailed in an article published in *Journal of the American Chemical Society*, by simulating the non-linear electronic response of DNA to irradiating protons and alpha-particles, the researchers examined the validity of the



A snapshot from a non-equilibrium electron dynamics simulation of solvated DNA under proton irradiation. Image: Dillon C. Yost, University of North Carolina at Chapel Hill

commonly used linear response theory description. The work found that the maxima for stopping power and electron-hole pair generation do not occur at the same irradiating ion velocity. A work published in *Physical Review Letters* examined the role of core-electron excitations in the electronic stopping of liquid water. K-shell core-electron excitations were found not only to provide an additional channel for the energy transfer but also to significantly influence the valence electron excitations. In the excitation process, generated holes remain highly localized within a few angstroms of the irradiating proton paths, whereas electrons are excited away. The core electrons only minorly contribute to the electron-hole pair density generated. Building on these two publications, the team is currently analyzing the complicated simulations to answer how water solvation influences the electronic stopping of DNA under ion irradiation.

**IMPACT** Greater understanding of electronic stopping processes at the molecular scale will advance a variety of modern technologies and medicine, including proton-beam cancer therapy.

## PUBLICATIONS

Yao, Y., D. C. Yost, and Y. Kanai. "K-shell Core-Electron Excitations in Electronic Stopping of Protons in Water from First Principles," *Physical Review Letters* (August 2019), APS.

Yost, D. C., and Y. Kanai. "Excitation Dynamics in DNA under Proton and  $\alpha$ -Particle Irradiation," *Journal of the American Chemical Society* (March 2019), American Chemical Society.



# Constructing and Navigating Polymorphic Landscapes of Molecular Crystals

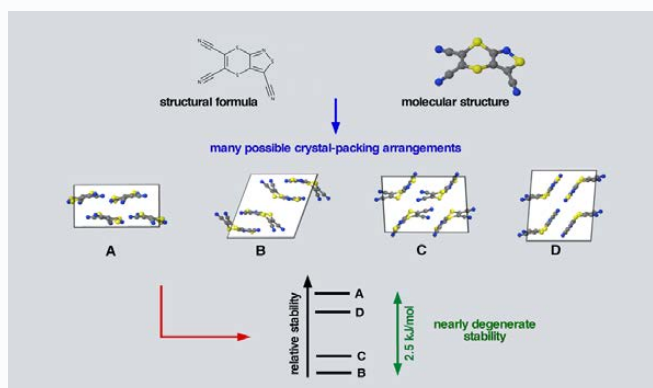
PI Alexandre Tkatchenko  
 INST University of Luxembourg  
 AWARD ADSP  
 HOURS Theta: 8,000,000 Node-Hours

**Molecular crystals are ubiquitous throughout science and technology, with numerous important applications. This project, led by University of Luxembourg researchers, seeks to combine state-of-the-art atomistic quantum simulations and data science methods to enable accurate predictions of novel molecular crystals for alternative energy materials, disease-curing pharmaceuticals, and molecular electronics.**

**CHALLENGE** Of crucial importance to the rational design and precise control over the physicochemical properties of a given molecular crystal is the fact that molecular solids frequently have many accessible polymorphs—that is, alternative structures that are nearly degenerate in stability yet display drastically different and potentially undesirable characteristics. Further progress in the discovery and understanding of composition–structure–property relationships of molecular crystals hinges on the availability of high-quality and systematic data-driven knowledge for structures, relative energies, and properties for many polymorphic systems. With this in mind, the primary goal of this research to combine an accurate, reliable, and computationally feasible quantum-mechanical methodology with cutting-edge data science techniques to enable first-principles-based predictive determination of the structure and stability of molecular crystal polymorphs.

**APPROACH** Within their quantum-mechanical and data-driven framework, the researchers carry out extensive state-of-the-art crystal structure prediction studies on ALCF supercomputers using a highly parallelized version of the density functional theory (DFT) code FHI-aims. All of the computed structures and energetics are to be made publicly available in a compiled database to guide further research and discoveries.

**RESULTS** After completing hybrid density functional theory calculations on 4.2 million structures of small-to-medium



Molecular crystal structure prediction for tricyano-1,4-dithiino[c]-isothiazole. The stablest molecular structure is calculated and used to create many possible crystal-packing arrangements. Image: Marcus A. Neumann, Avant-Garde Materials Simulation Deutschland GmbH; Robert A. Distasio Jr., Cornell University; Hsin-Yu Ko and Roberto Car, Princeton University; Johannes Hoja and Alexandre Tkatchenko, University of Luxembourg

molecules, the researchers compiled the results into a database with an initial analysis of molecular properties. The calculated reference data enable the development and testing of data-science approaches to describing the potential energy surfaces of individual molecules, which is a necessary prerequisite for the description of periodic systems like molecular crystals. The resulting database represents a new standard for molecular simulations; full analysis of its contents is expected to yield many breakthroughs.

**IMPACT** This work will help drive predictions of novel molecular crystals with numerous applications ranging from organic photovoltaics, molecular electronics, hydrogen sources, and pharmaceuticals, to agrochemicals, dyes, and food science. The database it yields will be an invaluable resource to the greater research and high-performance computing communities, enabling breakthroughs in computational materials discovery.

## PUBLICATIONS

Hoja, J., H.-Y. Ko, M. A. Neumann, R. Car, R. A. DiStasio Jr., and A. Tkatchenko. "Reliable and Practical Computational Description of Molecular Polymorphs," *Science Advances*, (January 2019), AAAS.



# Data-Driven Materials Discovery for Optoelectronic Applications

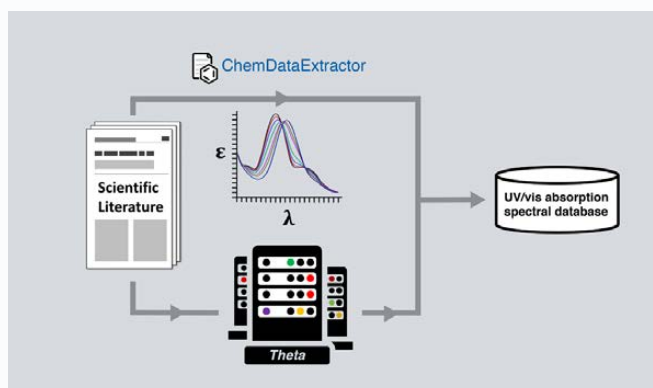
PI Jacqueline Cole  
INST University of Cambridge  
AWARD ADSP  
HOURS Theta: 468,750 Node-Hours

**Optoelectronics** is a branch of technology involving electronic devices that source, detect, and control light. Advances in this field are critical to many industries, including solar energy, data storage, and telecommunications. With this ADSP project, researchers from the University of Cambridge and Argonne National Laboratory are developing data-driven, materials-by-design capabilities to accelerate the discovery of new chemicals for optoelectronic applications.

**CHALLENGE** A majority of functional materials are discovered by trial-and-error methods. This unpredictable approach can present a bottleneck to technological innovation. To overcome this issue, the ADSP team is using systematic molecular design and engineering strategies to develop algorithms that mine massive chemical datasets to discover optimal materials for targeted applications.

**APPROACH** The workflow for this project involves four key steps: data extraction, data enrichment, materials prediction, and experimental validation. Using their text-mining tools, they extract data from academic literature to build a comprehensive repository of optoelectronic chemical and properties. To enrich the dataset, the team uses ALCF supercomputing resources to perform high-throughput electronic structure calculations, which produce paired quantities of experimental and computational data. The researchers then employ analytical methods to determine patterns in data that will lead to the development of algorithms and workflows that can predict structure-function relationships for new materials. Finally, the team carries out experimental work to validate the candidate materials.

**RESULTS** As part of a new study, the researchers are developing a pipeline to auto-generate an experimental and computational materials database of ultraviolet-visible (UV-vis) absorption spectra. They used their text-mining tool,



Auto-generating a UV-vis absorption spectral database via a dual experimental and computational chemical data pathway using the Theta supercomputer. Image: Jacqueline Cole and Ulrich Mayer, University of Cambridge

ChemDataExtractor, on more than 400,000 scientific documents to identify 18,309 experimentally determined UV-vis absorption maxima for the database. High-throughput electronic-structure calculations were carried out to predict the absorption maxima and oscillation strengths for a subset of the validated materials. They found large numbers of paired quantities of experimental and computational physical properties, laying the path for reliable in silico calculations of additional optical properties.

**IMPACT** The team's use of data mining, in conjunction with large-scale simulations and experiments, offers a novel approach to accelerate the design and discovery of materials for optoelectronic applications (e.g., solar cells, lighting, and data storage technologies). In addition, the project's development of open-source databases and data-extraction software tools will help remove the hurdle of manual database creation, creating an opportunity to stimulate further technological innovations in materials discovery.



# Imaging and Controlling Elemental Contrast of Nanocluster in Intense X-ray Pulses

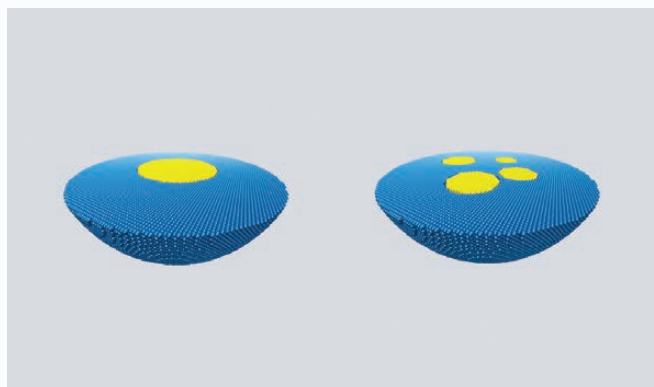
PI Phay Ho  
INST Argonne National Laboratory  
AWARD ALCC  
HOURS Mira: 5,625,000 Node-Hours

High-brightness x-ray free-electron laser (XFEL) pulses provide researchers with unprecedented tools to observe the dynamics of atoms and electrons with atomic spatial resolution. Researchers from Argonne National Laboratory aim to develop a quantitative and predictive understanding of x-ray matter interactions in nanosized heterogeneous systems on an atomistic level: Characterizing the complex phenomena that emerge can subsequently guide future experimental strategies.

**CHALLENGE** The outcome of high-intensity x-ray laser pulses with matter strongly depends on the pulse parameters and is preceded by a complex sequence of electronic transitions. By exploiting sample heterogeneity with these pulses, the researchers hope to control the ultrafast transient electronic and nuclear dynamics that occur and subsequently develop experimental schemes to enable imaging of complex materials with high spatial resolution and elemental contrast.

**APPROACH** The challenges associated with tracking the motion of particles and evolution of electronic configurations are addressed with a novel Monte-Carlo/molecular-dynamics (MC/MD) simulation algorithm implemented in the highly parallelized simulation code LAMMPS, consistent with the computational demands required to study multimillion-particle systems. At each step in the simulation, each particle coordinate is updated as a function of forces derived from interaction potentials. At each MC step, however, an MC procedure is used to incorporate quantum effects via cross sections and to evaluate probabilities for the occurrence of photoionization, resonant excitation, inner-shell relaxation collisional ionization, and electron-ion recombination.

**RESULTS** Carrying out x-ray/x-ray pump-probe calculations with heterogeneous nanoparticles, the researchers examined



Single-core vs many-core oxygen-doped Ne clusters. The yellow and blue particles depict oxygen and Ne atoms respectively. *Image: Phay Ho*

the dynamics and x-ray scattering images of two types of core-shell clusters with 30-nanometer diameters in intense XFEL pulses under different conditions. The first cluster consists of 8.2 million particles (electrons and nuclei), a single oxygen core of approximately 10 nanometers, and a neon shell approximately 10 nanometers thick. The second has the same number of oxygen and neon atoms but has multiple smaller oxygen clusters whose sizes range from 2 to 8 nanometers. These clusters consist of elements with similar atomic numbers, and serve as model biomolecules. The optimal pulse parameters that enable elemental contrast and highest spatial resolution are determined from the obtained scattering images.

**IMPACT** This work is essential in guiding future XFEL experiments with heterogeneous materials and establishing the applied methodology as an effective large-scale computational tool for the new research frontier of ultrafast x-ray science.



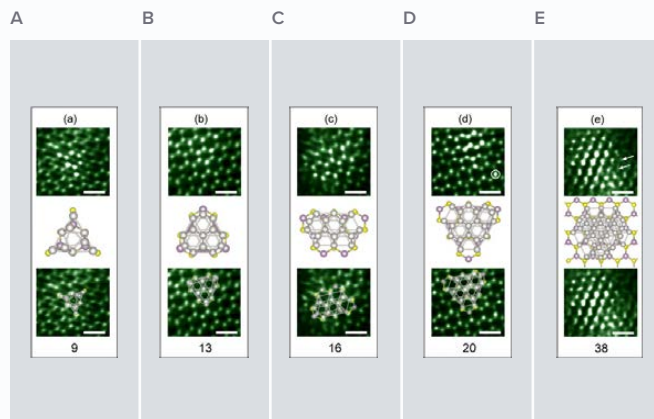
# Impact of Grain Boundary Defects on Hybrid Perovskite Solar Absorbers

PI Wissam A. Saidi  
 INST University of Pittsburgh  
 AWARD ALCC  
 HOURS Mira: 1,250,000 Node-Hours

**Environmental effects and intrinsic energy-loss processes lead to fluctuations in the operational temperature of solar cells, which can profoundly influence their power-conversion efficiency. Researchers from the University of Pittsburgh and Hamad Bin Khalifa University leveraged ALCC supercomputers to determine how the electronic structure of methylammonium lead iodide changes with temperature by accounting for electron-phonon coupling and thermal expansion.**

**CHALLENGE** Temperature effects can be modeled by accounting for the interactions between the electrons and phonons in the system. Because the electron-phonon coupling is challenging to compute using first-principles methods, the commonly employed Allen-Heine-Cardona theory resorts to a low-order expansion in terms of vibrational normal modes. However, for some systems, such as methylammonium lead iodide, the higher-order terms are important, and thus the low-order expansion is insufficient to capture the temperature dependence of the electronic structure. For these cases, a real-space approach that accounts for interactions to all orders is required, although this necessitates the use of massively parallel, leadership-class computing resources.

**APPROACH** The team computed the electron-phonon interactions for the orthorhombic and tetragonal phases of methylammonium lead iodide to all orders in electron-phonon coupling using a real-space Monte Carlo approach in conjunction with density functional theory calculations. These investigations took an approach similar the one the group had previously used on the cubic phase of methylammonium lead iodide with the smaller unit cell. Spin-orbit coupling is included, as this is found to enhance the electron-phonon coupling strength. On the other hand, nonlocal correlations are not accounted for, as these had little effect on the coupling.



Pseudo-color HAADF-STEM images of Pt clusters on MoS<sub>2</sub> nanoflake, DFT-calculated Pt structure, and the overlaid images. Cluster configurations for (A) 9 Pt, (B) 13 Pt, (C) 16 Pt, (D) 20 Pt and (E) 38 Pt. S, Mo, and Pt are shown as yellow, purple, and gray spheres. Image: Wissam A. Saidi, University of Pittsburgh; Samuel Poncé, University of Oxford; Bartomeu Monserrat, University of Cambridge; Ali Kachmar, Hamad Bin Khalifa University

**RESULTS** This project unambiguously confirmed for the first time the importance of high-order terms in the electron-phonon coupling by direct comparison with experiment. The researchers found that the band gap of methylammonium lead iodide increases with temperature, in excellent agreement with experimental results. They verified that anharmonic effects are only important near the tetragonal-cubic phase transition temperature. They also found that temperature has a significant effect on the effective masses and Rashba coupling in methylammonium lead iodide. At room temperature, electron-phonon coupling is found to enhance the band effective mass by a factor of two and to diminish the Rashba coupling by the same factor compared to T=0 K values. Numerous calculations performed for this work are detailed in a paper published in *The Journal of Physical Chemistry Letters*.

**IMPACT** This research, aligned with the DOE mission to advance clean energy, provides a comprehensive understanding of how temperature impacts the electronic structure of organic-inorganic perovskites that are of high interest for photovoltaic and optical applications.

## PUBLICATIONS

Saidi, W. A., and A. Kachmar. "Effects of Electron-Phonon Coupling on Electronic Properties of Methylammonium Lead Iodide Perovskites," *The Journal of Physical Chemistry Letters* (December 2018), ACS Publications.

Saidi, W. A., S. Poncé, and B. Monserrat. "Temperature Dependence of the Energy Levels of Methylammonium Lead Iodide Perovskite from First-Principles," *The Journal of Physical Chemistry Letters* (December 2016), ACS Publications.



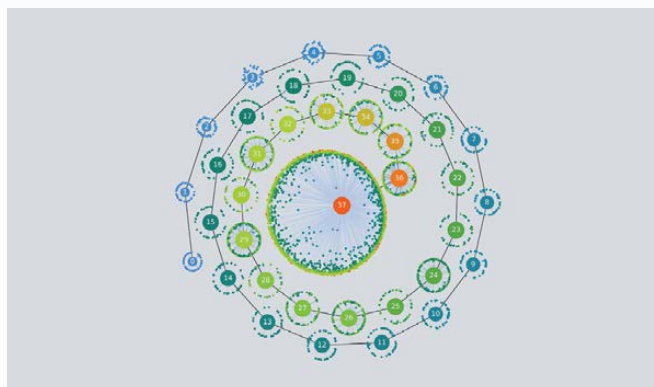
# Optimizing Superconductor Transport Properties through Large-Scale Simulation

PI Andreas Glatz  
INST Argonne National Laboratory  
AWARD DD  
HOURS Cooley: 12,000 Node-Hours

The ability of superconductors to carry large amounts of current at high magnetic fields is a key requirement for future design innovations in high-field magnets for accelerators and compact fusion reactors, and largely depends on the vortex pinning landscape comprised of material defects. With help from ALCF and OLCF computing resources, researchers from Argonne National Laboratory used the power of artificial intelligence to introduce and assess the impact of different configurations of defects on the performance of a superconductor.

**CHALLENGE** Too few defects cause a superconductor to become dissipative quickly, while too many defects can lead to blocking the super-current pathways or a breakdown of the superconducting material altogether. Therefore, scientists must be selective in how they incorporate defects into a material. With significant computing power and the development of sophisticated algorithms, researchers can systematically characterize and improve superconductors with different defects to find the optimal pinning landscape.

**APPROACH** To effectively tackle this problem with many degrees of freedom, the research team developed an AI approach based on genetic algorithms, which treat each defect like a biological gene. In contrast to conventional optimization techniques such as coordinate descent, where one varies only a few parameters characterizing the entire pinning landscape, this targeted evolution approach allows researchers to vary each defect individually without any a-priori assumptions about the defects configuration. The researchers leveraged ALCF computing resources to run the model. This effort is part of the OSCon (Optimizing SuperConductor Transport Properties through Large-Scale Simulation) partnership supported by DOE's Scientific Discovery Through Advanced Computing (SciDAC) program.



This image depicts the algorithmic evolution of a defect structure in a superconducting material. Each iteration serves as the basis for a new defect structure. Redder colors indicate a higher current-carrying capacity.  
Image: Andreas Glatz, Argonne National Laboratory

**RESULTS** To find the right combination of defects to arrest the motion of the vortices, the researchers initialized their algorithm with defects of random shape and size, giving the model a set of neutral initial conditions from which to work. As the researchers ran through successive generations of the model, they saw the initial defects transform into a columnar shape and ultimately a periodic arrangement of planar defects. The approach allowed the team to dynamically evolve the defect landscape into the best possible pinning configuration with maximal lossless current. A paper based on their study was published in the *Proceedings of the National Academy of Sciences*.

**IMPACT** The team's novel computational approach has the potential to advance the development of superconductors that can transmit electric current as efficiently as possible. In addition, their methodology to improve the intrinsic properties of condensed matter systems provides a promising path toward the design of tailored functional materials. It can be applied to a large variety of different physical systems and has demonstrated its usefulness in enhancing superconducting critical currents.

## PUBLICATIONS

Sadovskyy, I. A., A. E. Koshelev, W.-K. Kwok, U. Welp, and A. Glatz. "Targeted Evolution of Pinning Landscapes for Large Superconducting Critical Currents." *Proceedings of the National Academy of Sciences* (May 2019).



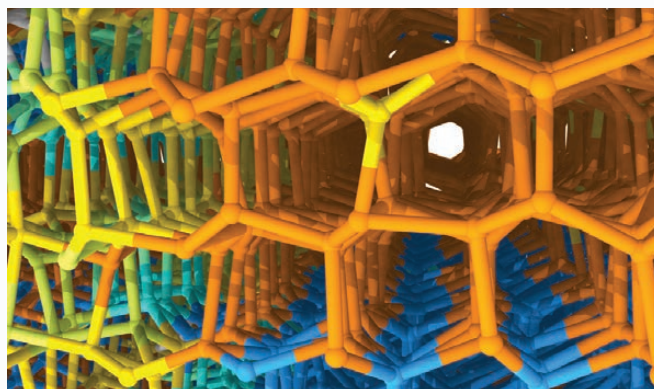
# Phase Transitions in Water-Ice-Vapor System

PI Subramanian Sankaranarayanan  
INST Argonne National Laboratory  
AWARD DD  
HOURS Mira: 1,160,000 Node-Hours

While water is perceived to be one of the simplest substances in the world, modeling its behavior on the molecular level is a complex and challenging task. Leveraging ALCF supercomputers, a team from Argonne National Laboratory employed machine learning to develop a new, computationally inexpensive water model that more accurately represents the thermodynamic properties of water, including how water changes to ice.

**CHALLENGE** Accurate and computationally efficient molecular-level descriptions of mesoscopic behavior of ice-water systems remain a challenge despite continued advances in computing hardware. For example, a molecular-level picture of phase transformations in water is most desirable, but remains inaccessible to fully atomistic simulations due to system size and timescale limitations. Making use of a hierarchical optimization workflow and leadership-class computing resources, the Argonne research team was able to create machine-learned water models that correctly predict multiple physical properties of ice-water systems at a fraction of the computational cost of the atomistic water models currently available.

**APPROACH** The researchers used coarse-grained modeling to treat individual water molecules as single particles in order to perform the simulation at a low computational cost, while machine learning helped ensure that their simpler model maintained the accuracy often sacrificed in these coarse-grained models. They utilized the ALCF's Mira system to perform simulations of up to 8 million water molecules to study the growth and formation of interfaces in polycrystalline ice. The machine-learned coarse-grained models were trained against first-principles calculations, experimental results, and temperature-dependent properties from other



Molecular dynamics simulations based on machine learning help scientists learn about the movement of the boundary between ice grains (yellow/green/cyan) and the stacking disorder that occurs when hexagonal (orange) and cubic (blue) pieces of ice freeze together. Image: Henry Chan and Subramanian Sankaranarayanan, Argonne National Laboratory

simulations using a multilevel, hierarchical global optimization strategy.

**RESULTS** The models were able to successfully predict both the experimental melting point of ice and the density maximum of liquid water at a computational cost two to six orders of magnitude lower than existing water models. The reduced computational cost of these models relative to their atomistic counterparts makes them suitable for modeling mesoscopic phenomena involving polycrystalline ice such as mechanics of ice, ice grain growth, melting of ice crystals, and pollutant effects on ice nucleation. The team's findings were published in *Nature Communications*.

**IMPACT** The team's machine learning-based model provides a tool that can accurately describe structural and thermodynamic properties of water and ice at a significantly lower computational cost than existing atomistic models. In addition, the machine learning workflow developed as part of this effort could be used to improve the performance and predictive capabilities of other atomistic and molecular models.

## PUBLICATIONS

Chan, H., M. J. Cherukara, B. Narayanan, T. D. Loeffler, C. Benmore, S. K. Gray, S. K. R. S. Sankaranarayanan. "Machine Learning Coarse-Grained Models for Water." *Nature Communications* (January 2019), Nature Publishing Group.



# Predictive Modeling and Machine Learning for Functional Nanoporous Materials

PI J. Ilja Siepmann  
 INST University of Minnesota  
 AWARD ALCC  
 HOURS Mira: 2,625,000 Node-Hours

The ability to identify optimal nanoporous materials for separation and catalytic processes has the potential to improve the production of biofuel and petroleum products, and advance the development of gas storage and carbon capture devices. An interdisciplinary team of researchers from the University of Minnesota is using predictive modeling and machine learning techniques on ALCF computing resources to accelerate the discovery and design of nanoporous materials for a variety of energy-related applications.

**CHALLENGE** More energy-efficient separation processes are needed to reduce the energy footprint of the chemical industries, but the modeling of complex sorption equilibria presents a bottleneck. Because of the lack of experimental data available for complex mixtures, computationally expensive molecular simulation has been the most used tool for predicting phase and sorption equilibrium properties. With this ALCC project, the team is employing deep neural networks as an efficient surrogate of molecular simulations, allowing for the prediction of properties at a fraction of the computational cost while the amount of data required to achieve accurate predictions is much less demanding.

**APPROACH** The researchers have developed a multitask deep neural network (DNN), called SorbNet, that can predict multicomponent adsorption isotherms over a wide range of temperatures and pressures. The modeling workflow combines molecular simulations with machine learning to predict the sorption behavior of complex chemical systems and deliver optimal desorption conditions.

**RESULTS** The team performed 131,072 simulations on ALCF's Mira supercomputer via an ensemble computing workflow to generate absorption data. SorbNet was trained on the



The team's SorbNet deep neural network can provide a shortcut for finding optimal conditions for a desorptive drying step. Image: Andrew Yangzesheng Sun, University of Minnesota

simulation data to predict equilibrium loadings as a function of thermodynamic state variables. They found that SorbNet accurately predicted the sorption loading results from simulation parameters, while also demonstrating an optimization application to find the optimal temperature to remove remaining solvent impurities during a desorptive drying step within an adsorption-drying-desorption process. In addition, SorbNet showed the ability to utilize transfer learning to tackle different product compounds, solvents, and porous materials. The team detailed their work with SorbNet in a paper published in *Chemical Science*.

**IMPACT** The team's SorbNet tool can be used to optimize the process conditions for more energy-efficient separations, providing a means to significantly reduce energy consumption in the chemical, biorenewable, and petrochemical industries. More broadly, this research provides a new avenue for applying machine learning and artificial intelligence in conjunction with molecular simulations to advance computational materials discovery for a wide variety of applications.

## PUBLICATIONS

Sun, Y., R. F. DeJaco, and J. I. Siepmann. "Deep Neural Network Learning of Complex Binary Sorption Equilibria from Molecular Simulation Data," *Chemical Science* (2019).



# Ab-initio Nuclear Structure and Nuclear Reactions

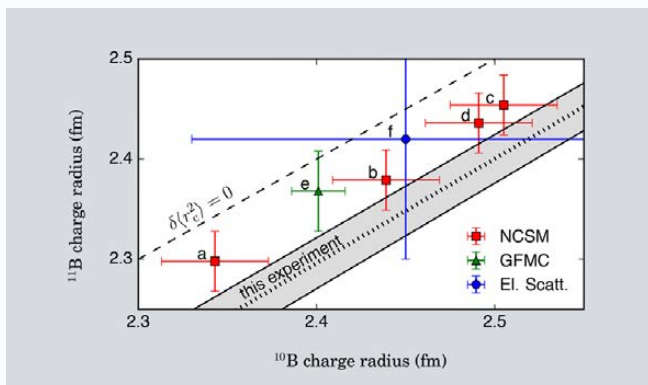
PI Gaute Hagen  
 INST Oak Ridge National Laboratory  
 AWARD INCITE  
 HOURS Mira: 3,500,000 Node-Hours;  
 Theta: 750,000 Node-Hours

Atomic nuclei are strongly interacting, self-bound many-body systems of fermions that display fascinating quantum mechanical properties. This project, led by researchers from Oak Ridge National Laboratory, builds on previous INCITE research to focus on first-principles approaches to nuclear structure and reactions that apply interactions derived within effective theories of quantum chromodynamics. The calculations performed will make predictions for and guide new experiments at major DOE facilities, explain observed phenomena, and potentially propel the discovery of new laws of nature.

**CHALLENGE** This project employs advanced quantum many-body methods to accurately describe and predict properties of the atomic nucleus from first principles, including their electroweak transitions and reactions important to both terrestrial experiments and astrophysical environments.

**APPROACH** The researchers leverage ALCF and OLCF supercomputing systems to employ advanced ab-initio quantum many-body techniques coupled with applied mathematics and computer science methods. This setup is used to perform ab-initio calculations of light- to heavy-mass nuclei and the equation of state of nuclear matter starting from the individual two- and three-nucleon interactions among their constituents: protons and neutrons. State-of-the-art simulations provide quantified simulations where direct experiment is impossible or subject to large uncertainties.

**RESULTS** As detailed in a paper published in *Physical Review Letters*, the researchers successfully modeled the stable isotopes boron-10 and boron-11 on Mira. Their theoretical calculations for nuclear charge radii are consistent with experimental results (published jointly) obtained via laser spectroscopy. Examination of the



Theoretical nuclear charge radii as obtained from Green's function Monte Carlo calculations performed on Mira and no-core shell model calculations with different nuclear potentials. The grey band represents the difference between  $^{11}\text{B}$  and  $^{10}\text{B}$  radii as measured by laser spectroscopy. Image: Bernhard Maaß, Institut für Kernphysik, Technische Universität Darmstadt

isotopes' nuclear configurations led to the determination that the 11 nucleons in boron-11 counterintuitively occupy a smaller volume than the 10 in boron-10.

**IMPACT** The calculations performed are relevant to numerous applications in nuclear energy, nuclear security, and nuclear astrophysics. The predictions obtained complement DOE investments in forefront experimental facilities and neutrino experiments. In this particular example, the agreement between theory and experiment helps validate predictions that the radioactive nucleus boron-8, for which an upcoming experiment at Argonne's ATLAS facility is planned, will occupy a larger volume than either boron-10 or boron-11.

## PUBLICATIONS

Maaß, B., T. Hünther, J. Krause, J. Krämer, K. König, A. Lovato, P. Müller, M. Puchalski, K. Pachucki, R. Roth, R. Sánchez, F. Sommer, R. B. Wiringa, and W. Nörtershäuser. "Nuclear Charge Radii of  $^{10,11}\text{B}$ ," *Physical Review Letters* (May 2019), American Physical Society.



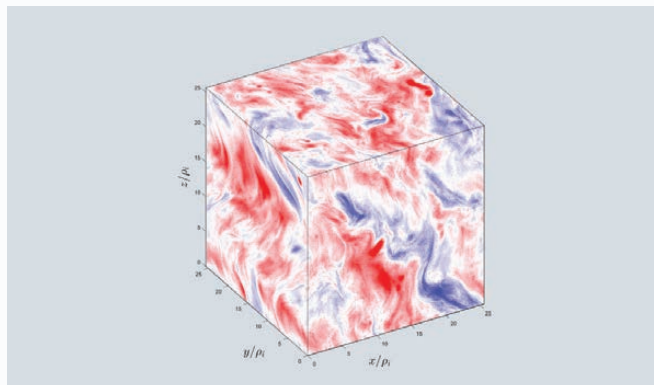
# Astrophysical Particle Accelerators: Magnetic Reconnection and Turbulence

PI Dmitri Uzdensky  
 INST University of Colorado Boulder  
 AWARD INCITE  
 HOURS Mira: 6,750,000 Node-Hours

Magnetic reconnection and turbulence play significant roles in a wide variety of plasma environments, from tokamaks and Earth's magnetosphere, to the solar corona and distant astrophysical sources where they can power intense gamma-ray flares. They are both fundamental plasma processes with important consequence, including the conversion of electromagnetic energy to particle kinetic energy. Besides heating the plasma, these processes can drive nonthermal particle acceleration (NTPA), resulting in a power-law particle energy distribution extending to very high energies. Understanding these mechanisms in the context of relativistic collisionless plasmas is especially important for high-energy astrophysical sources, where the existence of nonthermal particle populations is inferred from observations of nonthermal power-law radiation spectra. The acceleration mechanism, however, is a subject of continuing debate. With this project, researchers from the University of Colorado Boulder are studying relativistic particle energization and resulting radiative signatures in reconnecting and turbulent plasmas, with the ultimate goal of determining the conditions and mechanisms operating in astrophysical sources.

**CHALLENGE** Current research being performed for this project extends previous work on 3D collisionless magnetic reconnection and kinetic turbulence from electron-positron plasmas to electron-proton plasmas.

**APPROACH** The team is using its open-source particle-in-cell (PIC) code Zeltron, which is capable of self-consistently incorporating the radiation reaction force to model cases where radiative losses significantly alter system dynamics and NTPA. The massively parallel, 3D simulations of reconnection and turbulence permit the researchers to examine the viability of these mechanisms as potential efficient astrophysical



Surface image of the electron energy density from a snapshot of a 10243-cell particle-in-cell simulation of driven kinetic turbulence in a plasma consisting of sub-relativistic ions and ultra-relativistic electrons. Image: Vladimir Zhdankin, Princeton University

accelerators of large numbers of particles to ultrarelativistic energies sufficient to explain observed radiation.

**RESULTS** Through one extremely large simulation and a host of moderately large ones, the researchers were able to correlate with the formation time of the electron energy spectrum a previously discovered system-size dependence of NTPA. Further leveraging these simulations, the researchers are delving deeper into the microphysical mechanisms behind turbulent particle acceleration.

**IMPACT** These results will lead to advances in the understanding of fundamental plasma physics processes and have important implications for modern high-energy astrophysics.

## PUBLICATIONS

Zhdankin, V., D. A. Uzdensky, G. R. Werner, and M. C. Begelman. "Electron and Ion Energization in Relativistic Plasma Turbulence," *Physical Review Letters* (February 2019), American Physical Society.



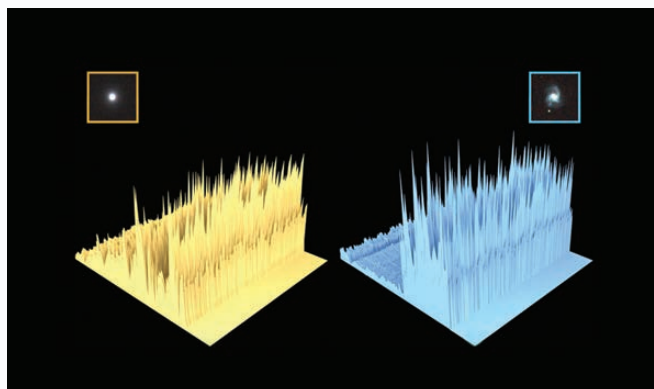
# Deep Learning at Scale for Multi-Messenger Astrophysics

PI Eliu Huerta  
 INST National Center for Supercomputing Applications and University of Illinois at Urbana, Champaign  
 AWARD ADSP  
 HOURS Theta: 156,250 Node-Hours

**Multi-messenger astrophysics (MMA)** refers to the observation of astrophysical phenomena using a multitude of cosmic messengers, including gravitational waves, electromagnetic waves, neutrinos, and cosmic rays. The astronomical facilities that probe the cosmos using these cosmic messengers produce datasets whose volume and velocity challenge the scalability and computational performance of standard data analytics algorithms. Researchers from the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign are spearheading a science program at the interface of AI and HPC to address these challenges to maximize discovery with MMA discovery campaigns in the big-data era.

**CHALLENGE** In collaboration with researchers from Argonne National Laboratory, this project is developing AI algorithms to significantly increase the depth and speed of gravitational wave searches, and to process in real-time terabyte-size datasets of telescope images to identify the electromagnetic counterparts of gravitational wave sources.

**APPROACH** Leveraging ALCF supercomputing resources, the researchers are using TensorFlow and PyTorch to design neural network models for real-time detection and characterization of gravitational wave sources, as well as the classification of galaxy images in two electromagnetic surveys, namely the Sloan Digital Sky Survey (SDSS) and the Dark Energy Survey (DES). The main goal is to use gravitational wave observations of compact binary mergers to establish a redshift-distance relationship that may be used to measure the Hubble constant. Working together with ALCF staff, this group has deployed optimized neural network models to be trained at scale with the open-source deep learning framework Horovod.



This image is a snapshot from a visualization that shows the output of the penultimate layer of a deep neural network during training as it learns to classify spiral and elliptical galaxies.. Image: Janet Knowles, Joseph A. Insley, and Silvio Rizzi, Argonne National Laboratory

**RESULTS** A paper published in *Physical Review D* demonstrates that deep learning can extract the complex gravitational wave signatures of eccentric binary black hole mergers from LIGO data. As shown in work in *Physical Letters B*, the team achieved the first application of deep transfer learning and distributed training to design neural network models for classifying galaxies in the SDSS and DES with state-of-the-art accuracy. This study also introduced a method to classify unlabeled galaxies in the DES. Other work demonstrated that deep learning algorithms can denoise gravitational wave signals embedded in non-Gaussian and non-stationary advanced LIGO noise. These new methodologies outperform traditional signal-processing techniques.

**IMPACT** Future MMA observations will provide information about the (currently unknown) matter composition of neutron stars, the origin of long gamma ray bursts, and the mass distribution of neutron stars. They also stand to elucidate the mechanisms that may trigger electromagnetic radiation when black holes and neutron stars collide, and will enable us to pinpoint the location of these events with unprecedented accuracy. Furthermore, the algorithms and computational frameworks designed and deployed in this project will benefit domain scientists and help promote deep learning as the default signal-processing method for big-data analytics.

## PUBLICATIONS

Khan, A., E. A. Huerta, S. Wang, R. Gruendl, E. Jennings, and H. Zheng. "Deep Learning at Scale for the Construction of Galaxy Catalogs in the Dark Energy Survey," *Physical Letters B* (August 2019), Elsevier.

Rebei, A., E. A. Huerta, S. Wang, S. Habib, R. Haas, D. Johnson, and D. George. "Fusing Numerical Relativity and Deep Learning to Detect Higher-Order Multipole Waveforms from Eccentric Binary Black Hole Mergers," *Physical Review D* (August 2019), American Physical Society.



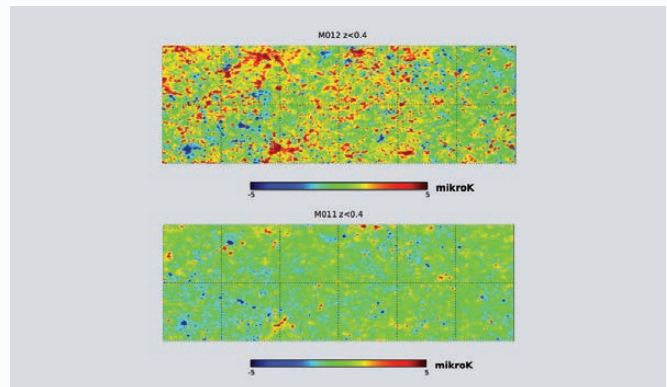
# Emulating the Universe

PI Katrin Heitmann  
 INST Argonne National Laboratory  
 AWARD ALCC  
 HOURS Mira: 625,000 Node-Hours

With this project, Argonne National Laboratory researchers aim to create a rich set of cosmological simulations that span different cosmological models to further our understanding of the make-up and history of the universe. Each cosmological simulation in the finalized Mira-Titan Universe suite will be transformed into a synthetic sky that closely mimics actual observational data, allowing for a meaningful comparison with observations, including a rigorous treatment of systematic errors. The project will use the simulation suite to build a set of high-accuracy prediction tools, so-called emulators. Emulators play an important role in the analysis of cosmological measurements by allowing the exploitation of the nonlinear regime of structure formation which holds important clues about the reason for the accelerated expansion of the universe.

**CHALLENGE** The unveiling of the large-scale structure of the universe by modern sky surveys has transformed cosmology, fundamental physics, and astrophysics. The resultant Standard Model is a remarkable success, but it raises deep questions about the nature and origin of its fundamental constituents and assumptions, such as dark energy, dark matter, and primordial fluctuations. Next-generation observatories are working to answer these questions and make discoveries beyond the Standard Model, but their promise can only be realized via multi-wavelength simulations, which play a crucial role in interpreting observations and providing predictions for cosmological models. However, such simulations carry steep computational requirements, necessitating use of leadership-class computing systems.

**APPROACH** The simulations are carried out with the Hardware/Hybrid Accelerated Cosmology Code, HACC, developed to execute extremely large, gravity-only cosmological simulations. HACC was designed to take full



Maps showing the predicted kSZ effect for two different cosmological models from the Mira Titan Universe suite. While the overall structures are in the same places it is clear that one of the models has much more structure than the other. Image: HACC team, Argonne National Laboratory

advantage of contemporary and future hardware architectures and has been optimized to run on Mira and the OLCF's Titan system. The analysis toolkit CosmoTools is used to extract survey-relevant science from the simulations, from which multi-wavelength synthetic maps are built (including models of the kinematic and thermal Sunyaev-Zel'dovich effect, weak lensing maps, and optical catalogs). A data portal utilizing the data management and sharing project Petrel has also been constructed so as to grant access to the greater cosmological community.

**RESULTS** The researchers carried out analysis tasks on the Mira-Titan universe simulations to enable the creation of a set of emulators. These emulators generating new power spectra. The researchers tested the emulator predictions against a set of five new simulations carried out as part of this project.

**IMPACT** The project aims to build next-generation computational cosmology prediction and analysis frameworks that are directly targeted to current and future observations. The simulations will provide important testbeds for new developments in building complex, yet robust, cosmological prediction tools.



Physics | 🧑🔬 Simulation

# Extreme-Scale Simulation of Supernovae and Magnetars from Realistic Progenitors

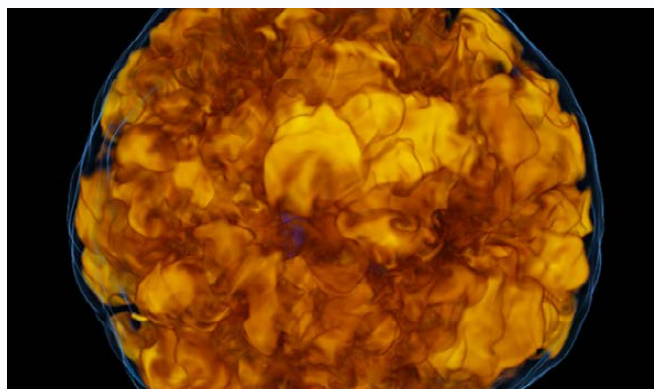
PI Sean Couch  
INST Michigan State University  
AWARD INCITE  
HOURS Mira: 9,375,000 Node-Hours;  
Theta: 281,000 Node-Hours

**Core-collapse supernovae (CCSNe) produce violent explosions that mark the deaths of massive stars. Key to cosmic chemical evolution, these events synthesize most of the elements heavier than helium throughout the universe. Despite the critical role supernovae play in astrophysics, the physical mechanism that causes the stellar explosions is still not fully understood. A team from Michigan State University is using ALCF computing resources to perform extreme-scale, 3D simulations aimed at transforming our understanding of supernovae.**

**CHALLENGE** Exploration of this complex multiphysics problem, especially via high-fidelity, fully 3D simulations, is an extremely computationally demanding task. Recent advances in computing hardware and software, however, have made it possible for researchers to employ 3D simulations that can both capture crucial hydrodynamic instabilities and model multidimensional neutrino transport. This emerging capability has driven significant advances in the theoretical study of CCSNe.

**APPROACH** The research team's comprehensive, end-to-end investigation involves carrying out 3D magnetohydrodynamics (MHD) simulations with sophisticated multidimensional neutrino transport and the most realistic initial conditions ever adopted to study the effects of rotation, magnetic fields, and progenitor asphericity on the CCSNe mechanism. Their research will also enable an intensive comparison to observations through the calculation of gravitational wave emission, detailed nucleosynthesis, and electromagnetic radiative transfer. The primary tool for the team's work at the ALCF is the FLASH hydrodynamics code, which implements a state-of-the-art, high-order MHD solver in an adaptive mesh refinement framework.

**RESULTS** Using simulations on Mira to examine various In a study recently published in the *Astrophysical Journal*, the



Volume rendering of entropy from a 3D simulation of a magnetorotational core-collapse supernova simulation carried out on Mira. This simulation uses a realistic rotation profile and speed taken from a stellar evolution model that included the effects of rotation and magnetic fields. Image: Sean Couch, Michigan State University

team detailed the results of an extensive set of 3D CCSNe simulations using high-fidelity multidimensional neutrino transport, high-resolution hydrodynamics, and approximate general relativistic gravity. Their simulations showed that increased turbulent activity can play a crucial role in supernova shock dynamics by providing additional pressure support and increased neutron heating, quantitatively bringing CCSNe closer to the point of explosion.

The team is currently using ALCF computing resources to perform simulations aimed at predicting gravitational wave emissions from supernovae.

**IMPACT** This project is performing massive 3D simulations to determine the impact of rotation and magnetic fields on the supernova explosion mechanism. The team's findings will inform our understanding of the characteristics of newborn pulsars and magnetars—information that can be directly compared to observational data. A better understanding of supernovae will also shed new light on the chemical evolution of the universe.

**PUBLICATIONS**  
O'Connor, E. P., S. Couch. "Exploring Fundamentally Three-Dimensional Phenomena in High-Fidelity Simulations of Core-Collapse Supernovae," *The Astrophysical Journal* (September 2018).



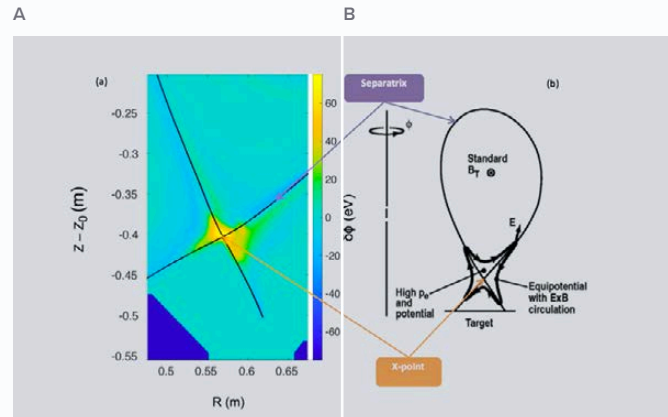
# High-Fidelity Gyrokinetic Simulation of Tokamak and ITER Edge Physics

PI Choong-Seock Chang  
 INST Princeton Plasma Physics Laboratory  
 AWARD INCITE  
 HOURS Theta: 1,500,000 Node-Hours

ITER, an international collaboration to design, construct, and assemble a burning plasma experiment that can demonstrate the scientific and technological feasibility of fusion energy, is a massive effort that relies on contributions from physicists around the world. This includes a team from the Princeton Plasma Physics Laboratory that is leading a three-year INCITE project aimed at furthering the understanding of the edge plasma physics in fusion reactors.

**CHALLENGE** There are two high-priority challenges in the advancement of ITER's quest toward fusion energy: (1) the peak heat-flux density on the ITER divertor target plates in the high-confinement mode (H-mode) operation and (2) the achievability of the H-mode plasma condition (i.e., achieving the low-to-high confinement mode L-H transition at edge) that is essential for ITER to meet its tenfold energy gain goal. Understanding these problems is an over 35-year-old issue, difficult due to the complexity of the tokamak edge geometry and the nonlinear, multiscale, multiphysics nature of the problems.

**APPROACH** To address these challenging plasma physics questions, the PPPL team is using its 5D gyrokinetic particle-in-cell code, XGC1, to perform large-scale simulations on DOE leadership computing resources, including the ALCF's Theta system. XGC1 has the unique capability to solve the boundary multiscale plasma problems across the magnetic separatrix (the boundary between the magnetically confined and unconfined plasmas) using first-principles-based kinetic equations. The magnetic field lines in the separatrix run through the X-point and continue outward, allowing waste heat and particles to escape and hit the divertor plate. If the energy deposited on the plate is too large or concentrated in too narrow an area, it may physically damage it. By



(A) Electrostatic potential in XGC1 simulation of H-mode in C-Mod tokamak, showing ExB flow of plasma around the magnetic X-point. (B) ExB flow near X-point from DIII-D tokamak measurements (in agreement with simulation). Image: Seunghoe Ku, Princeton Plasma Physics Laboratory

running XGC1 on Theta, the researchers are able to study the heat exhaust produced by fusion reactors.

**RESULTS** The team's simulations have revealed a surprising effect of the X-point on the heat exhaust: a bump of electric field occurs at the X-point, causing plasma to circulate around it and preventing plasma particles from traveling directly between the upstream and downstream areas. This shows that the present assumption used for the interaction between these edge-plasma components was inadequate, and more accurate equations and improved reduced models can now be developed. The team's findings were published in *Physics of Plasma*.

**IMPACT** Fusion energy has the potential to provide a safe, abundant, and carbon-free source of power. Achieving a tenfold energy gain in ITER experiments relies on a deeper understanding of edge plasma physics through large-scale, high-fidelity simulation. DOE leadership computing resources, along with the team's advanced XGC1 code, can allow for more accurate predictions of escaping waste heat and particles and will help inform the development of future, large-scale fusion experiments that are less vulnerable to internal damage.

## PUBLICATIONS

Chang, C. S., S. Ku, R. M. Churchill. "X-point Ion Orbit Physics in Scrape-off Layer and Generation of a Localized Electrostatic Potential Perturbation Around X-point," *Physics of Plasma* (January 2019), American Institute of Physics.



# Kinetic Simulation of FRC Stability and Transport

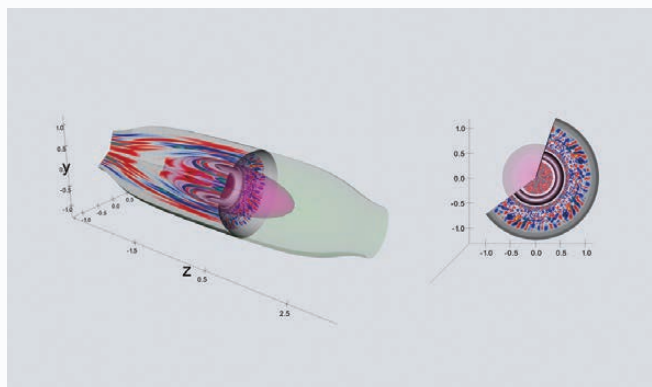
PI Sean Dettrick  
 INST TAE Technologies, Inc.  
 AWARD INCITE  
 HOURS Theta: 750,000 Node-Hours

**Fusion energy offers the prospect of a carbon-neutral, environmentally responsible, inexhaustible energy source. Researchers from TAE Technologies, Inc. are using ALCF computing resources to accelerate their fusion plasma research program aimed at developing the world's first commercially viable fusion-powered generator for electricity production.**

**CHALLENGE** The central challenge in generating fusion energy for commercial use is to keep plasma hot enough for long enough to sustain the fusion reaction. TAE Technologies has built "Norman," an advanced beam-driven field-reversed configuration (FRC) plasma device, for experimental studies of the magnetic confinement of hot fusion plasmas. Heat can be lost from the FRC due to nonlinear wave-particle interactions, which start at the microscale and lead to kinetic turbulence. On the macroscale, nonlinear wave-particle interactions can also determine the global stability of the plasma.

**APPROACH** To understand both the microscale and macroscale kinetic plasma physics in FRC devices, the research team has utilized Theta to develop the ANC kinetic micro-turbulence code and the FPIC kinetic macro-stability code to enable massively parallel 3D particle-in-cell (PIC) simulations of the Norman experiment. The ANC code, developed in collaboration with the University of California, Irvine, is being used to interpret experimental measurements of turbulence and to understand how heat loss (transport) scales with plasma temperature. The FPIC code is being used to help understand the global stability and self-organization of the plasma.

**RESULTS** Developing these codes has allowed the TAE Technologies team to increase physics understanding by comparing experimentally measured and computationally simulated FRC turbulence for the first time. Using the ANC code on Theta, the researchers performed the first non-linear,



3D simulations of micro-turbulence in the field-reversed configuration performed using the ANC code on Theta. Electrostatic fluctuation after saturation is shown from two different perspectives. Image: Daniel Fulton, Kevin Hubbard, and Calvin Lau, TAE Technologies, Inc.

global, cross-separatrix, turbulence simulations in FRC topology. The study revealed that short wavelength modes are linearly unstable outside of the FRC, but not on the interior. The simulations were validated against previous experimental observations and were found to agree well. These findings were published in *Nuclear Fusion*.

The team is using the FPIC code to study the global configuration stability of the FRC, simulating the impact of external actuators on plasma self-organization. These results will be used to inform the operating states and the feedback and control strategies that are being developed for Norman and future FRC devices.

**IMPACT** This project will help accelerate TAE Technologies' studies of the confinement of energy with high plasma temperatures and inform the design of a future prototype reactor. If the company's technology can be extended to fusion-relevant temperatures, it will provide an economical path to clean, safe, and sustainable fusion energy.

## PUBLICATIONS

Lau, C. K., D. P. Fulton, J. Bao, Z. Lin, T. Tajima, L. Schmitz, S. Dettrick, and the TAE Team. "Cross-Separatrix Simulations of Turbulent Transport in the Field-Reversed Configuration." *Nuclear Fusion* (May 2019), IOP Publishing.



# PLASM-IN-SILICO: HPC Modeling of High-Intensity Laser-Solid Interaction

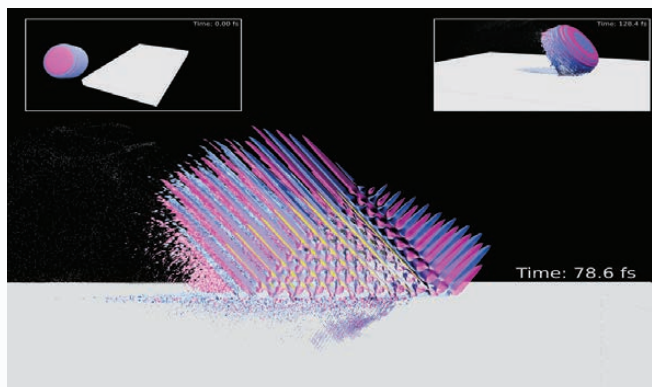
PI Jean-Luc Vay  
 INST Lawrence Berkeley National Laboratory  
 AWARD INCITE  
 HOURS Mira: 8,000,000 Node-Hours;  
 Theta: 680,000 Node-Hours

The advent of petawatt (PW) lasers has made it possible to achieve light intensities at which matter turns into a plasma. This technology has laid the groundwork for the new, promising—but still largely unexplored—field of ultra-high-intensity (UHI) physics, which investigates light-matter interactions at extreme intensities to understand the complex laws governing plasma dynamics in ultra-relativistic regimes. This project, led by researchers from Lawrence Berkeley National Laboratory in the US and Commissariat à l'Energie Atomique in France, uses massively parallel particle-in-cell (PIC) simulations on ALCF supercomputers to overcome challenges facing UHI physics.

**CHALLENGE** The researchers intend to develop novel PW laser-based solutions for three major outstanding problems in UHI: (1) Can we produce high-charge compact electron accelerators with high-beam quality that will be essential to push forward the horizons of high-energy science? (2) Can we produce efficient and compact high-energy ion accelerators to democratize cancer hadron-therapy? (3) Can we reach extreme light intensities approaching the Schwinger limit, beyond which light in a vacuum self-focuses and electron-positron pairs are produced?

In particular, this project aims to show that so-called “relativistic plasma mirrors,” produced when a high-power laser hits a solid target, can provide simple and elegant paths to solving these three challenges.

**APPROACH** The researchers are devising solutions in silico with massively parallel simulations on the ALCF’s Mira and Theta systems using the Warp and WarpX codes. To this end, their work leverages recent transformative developments in the first-principles simulation of UHI laser-plasma interactions that enabled the 3D, high-fidelity modeling of plasma mirror



Interaction between an ultra-high intensity laser (dark blue grey/magenta) and a plasma mirror (grey). Image: Guillaume Blaclard, CEA Saclay and Lawrence Berkeley National Laboratory

sources. Close collaborations with teams at PW laser facilities help enable experimental validation of the devised solutions.

**RESULTS** A paper recently accepted for publication in *Physical Review Letters* used cutting-edge 3D PIC simulations to demonstrate that relativistic plasma mirrors irradiated by PW lasers and naturally curved by laser radiation pressure can be used to tightly focus Doppler-generated harmonics to extreme intensities. The 3D simulations were employed to develop and validate a general 3D model of harmonic focusing by a curved relativistic plasma mirror. This model was used to propose novel all-optical techniques to further increase plasma mirror curvature and approach the Schwinger limit.

**IMPACT** Successful identification of optimal laser-plasma parameters for ion generation will enable experimental facilities to generate high-energy ion beams for numerous applications, including radiography, fast ignition, warm dense matter studies, nuclear physics, radiation effects in electronics, and hadron therapy. Devising novel schemes to reach extreme intensities should pave the way to the study of yet unexplored regimes of non-linear QED.

## PUBLICATIONS

Vincenti, H. “Achieving Extreme Light Intensities Using Optically-Curved Relativistic Plasma Mirrors,” *Physical Review Letters* (September 2019), American Physical Society.



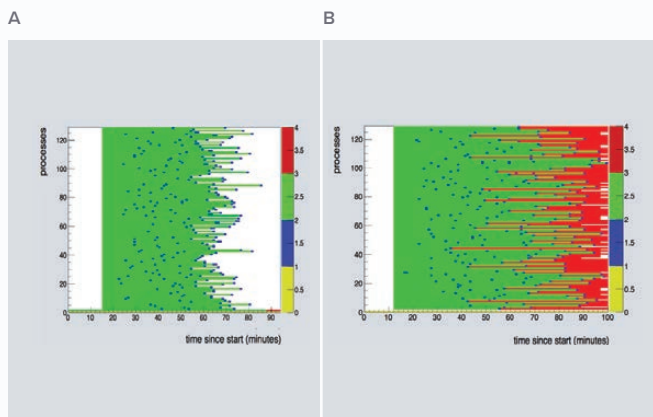
# Scaling LHC Proton-Proton Collision Simulations in the ATLAS Detector

PI Eric Lancon  
 INST Brookhaven National Laboratory  
 AWARD ALCC  
 HOURS Theta: 1,250,000 Node-Hours

**With the High-Luminosity Large Hadron Collider (LHC) looming on the horizon, the data and computing challenges facing the LHC physics program will increase by more than an order of magnitude. With this project, researchers led by Brookhaven National Laboratory are continuing the ongoing effort within ATLAS to adapt and evolve its software for efficient use of DOE supercomputers, including the development of a natively multithreaded simulation framework, investigating the use of checkpointing, profiling file system operations, and running simulation workflow steps back-to-back to streamline data storage and transfer.**

**CHALLENGE** The Standard Model of particle physics is being tested with increasing precision in proton-proton collisions by the ATLAS experiment at the LHC. While the masterpiece of the model, the Higgs boson, has been discovered at the LHC, many fundamental questions have no answer within the present understanding of particle physics: Why is there more matter than antimatter? What is dark energy? Why do neutrinos have mass? Answering these questions requires not just precision experiments be performed at the LHC, but that those experiments be compared with the results of precision model simulations. This project harnesses ALCF supercomputers in order to drive software optimization to improve the computational efficiency of such models, and to provide 160 million simulated proton collisions.

**APPROACH** This work extends efforts to efficiently deploy Athena, the ATLAS simulation and analysis framework, on Theta and to evolve PanDA, a workflow management system. A modular edge service, Harvester, is being developed as a unified approach for harnessing DOE supercomputer to communicate with PanDA, stage input data, launch jobs on the local batch system, communicate with batch jobs to provide jobs, and upload output data.



One KNL node running AthenaMP on 128 KNL threads without (A) and with (B) the Event Service. KNL threads are represented by one thread per row and color indicates their activity. Green indicates the thread is processing events, white means the thread is idle, and red indicates events which were lost due to end-of-wall-time reached. Image: Taylor Childers, Argonne National Laboratory

**RESULTS** Utilizing the Argonne-developed Yoda software package on Theta, the ATLAS experiment ran fine-grain workflows to coordinate and orchestrate the ATLAS Monte Carlo simulation programs for modeling various processes in particle physics. The first quarter of 2019 saw ATLAS simulate over 90 million collisions on Theta.

**IMPACT** Mounting computational and data challenges make DOE supercomputers necessary for the LHC to meet its scientific goals. By supporting the ATLAS project within the DOE Office of Science, this work will guide both experimentalists and theorists searching for physics beyond the Standard Model, helping to advance our understanding of fundamental particles and forces.



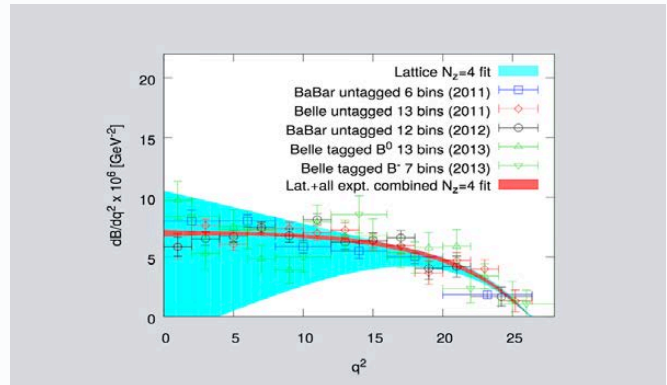
# Semileptonic B- and D-meson Form Factors with High Precision

PI Ruth Van de Water  
 INST Fermilab  
 AWARD ALCC  
 HOURS Mira: 1,250,000 Node-Hours;  
 Theta: 1,093,750 Node-Hours

Despite the enormous success with which the Standard Model of particle physics describes the subatomic world, several mysteries connected to quark interactions, neutrino physics, and cosmology cannot be reconciled without additional particles and/or new forces among known particles. One such mystery is the preponderance of matter over antimatter in the universe, which can currently only be explained by different interaction rates (a circumstance known as charge conjugation parity symmetry violation or CP violation). The CP-violating processes in the Standard Model aren't strong enough to explain our universe, meaning more such interactions must exist. One hypothesized source of new CP violation is quark interactions, but experiments cannot detect quarks, which are always bound into larger particles called hadrons. This project seeks to use lattice quantum chromodynamics (QCD) to compute the connection between quarks and hadrons with unsurpassed precision.

**CHALLENGE** Precise calculations of hadron properties are necessary to advance elementary particle physics, but their execution require substantial computational power. Harnessing the power of ALCF supercomputers, this project builds on previous work to enable experiments that could identify possible new sources of CP-violation. The lattice-QCD calculations performed will reduce uncertainty in all semileptonic decays with a B- or D-meson in the initial states and a pion or kaon in the final state.

**APPROACH** To compute transition matrix elements (form factors) for vector, scalar, and tensor currents, the researchers are using the MILC code, which has been optimized to run at scale on ALCF systems. There are two steps to lattice-QCD calculations: First, Markov-chain Monte Carlo generates



Partial branching fraction for the decay  $B \rightarrow D^* l \nu$ . The lattice-only fit result (cyan band) comes from our prior calculation of  $|V_{ub}|$  and the combined-fit result (red band) comes from fitting the experimental data together with the lattice form factor, leaving  $|V_{ub}|$  as an unknown parameter to match the two results. Image: Ruth Van de Water, Fermilab

ensembles of typical gauge-field configurations. Second, hadron-level quantities (such as form factors) are computed.

**RESULTS** Analyses have been completed for 1000 lattices, realizing the team's goal. Data from the analyses were used in a paper submitted to *Physical Review D* in order to calculate the contribution to the muon anomalous magnetic moment hadronic vacuum polarization from the connected diagrams of up and down quarks, omitting electromagnetism. Furthermore, calculations for the form factors of the  $B \rightarrow D^* l \nu$  semileptonic decay were reviewed, with lattice predictions compared with cutting-edge experimental results.

**IMPACT** The search for new particles and interactions lies at the heart of high-energy physics research. This work will be useful to groups working on quark-flavor physics and lattice QCD.

## PUBLICATIONS

Davies, C. T. H. et al. "Hadronic-vacuum-polarization contribution to the muon's anomalous magnetic moment from four-flavor lattice QCD," *Physical Review D* (submitted), American Physical Society.

Vaquero, A., C. Datta, A. X. El-Khadra, A. S. Kronfeld, J. Laiho, and R. S. Van de Water. "B to D\* l nu at Non-zero Recoil," *Proceedings of the 2019 Conference on Flavor Physics and CP Violation* (May 2019), Springer.



Physics |  Simulation

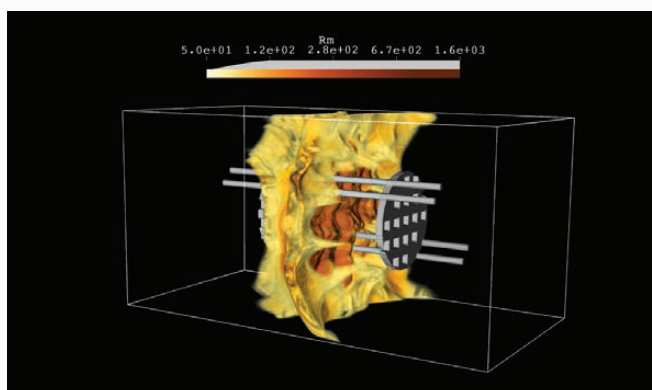
# Simulations of Laser Experiments to Study MHD Turbulence and Non-Thermal Charged Particles

PI Petros Tzeferacos  
 INST The University of Chicago and Argonne National Laboratory  
 AWARD ALCC  
 HOURS Mira: 1,375,000 Node-Hours

Magnetic fields are ubiquitous in the universe. The standard theoretical model of their origin is the amplification of tiny seed fields via turbulent dynamo, resulting in turbulent magnetized plasmas. These turbulent plasmas mediate the propagation and acceleration of non-thermal charged particles via scattering and second-order Fermi acceleration—a stochastic process whereby cosmic rays gain energy through scattering with randomly moving magnetized clouds. With this project, University of Chicago and Argonne National Laboratory researchers are using simulations to design, execute, and interpret laser-driven experiments that study the properties of magnetized turbulent plasmas and non-thermal particles in the laboratory.

**CHALLENGE** The turbulent dynamo mechanism eluded experimental demonstration for decades, and studies largely relied on simplified models that do not capture realistic magnetized turbulence as it occurs in astrophysical systems. Similarly, the Fermi mechanism still awaits experimental demonstration in a controlled environment, and its study has required simplified assumptions. The project has been awarded time at experimental facilities to demonstrate and study these processes in the laboratory, where extremely powerful lasers produce magnetized turbulence at the large magnetic Reynolds numbers ( $R_m$ ) required for the nonlinear turbulent dynamo mechanism to occur. The experiments are designed and interpreted using validated simulations whose execution necessitates leadership-class computers due to their complexity and the short time permitted by the scheduled experiments.

**APPROACH** The researchers are using Mira to run FLASH, a highly capable, parallel, multiphysics, adaptive mesh refinement, finite-volume Eulerian hydrodynamics and magnetohydrodynamics (MHD) code. The researchers used FLASH to design a platform that could generate



FLASH simulation of the Omega experiments. Shown is the volume rendering of the magnetic Reynolds number,  $R_m$ , demonstrating values well above the threshold for turbulent dynamo to operate. *Image: Petros Tzeferacos, The University of Chicago and Argonne National Laboratory*

high- $R_m$  magnetized turbulence and capture the kinematic, non-linear, and saturation phases of the dynamo.

**RESULTS** The FLASH simulation campaigns enabled breakthrough experiments at Omega and NIF, which demonstrated subsonic turbulent dynamo in the laboratory and Markovian diffusion of high-energy charged particles in spatially intermittent turbulent magnetic fields, creating an experimental analogue for cosmic-ray transport in the intergalactic medium. Via simulation, the researchers redesigned their experimental platform to demonstrate dynamo in the supersonic regime and, for the first time, capture second-order Fermi acceleration.

**IMPACT** This research establishes a basis for laboratory investigations into the nature of turbulent dynamo and MHD turbulence, furthers the transformation of the academic community's ability to design and analyze high energy density physics (HEDP) experiments at large laser facilities, and trains young scientists to design and interpret HEDP experiments using validated simulations.

## PUBLICATIONS

Tzeferacos, P., A. Rigby, A. F. A. Bott, A. R. Bell, R. Bingham, A. Casner, F. Cattaneo, E. M. Churazov, J. Emig, F. Fiuza, C. B. Forest, J. Foster, C. Graziani, J. Katz, M. Koenig, C.-K. Li, J. Meinecke, R. Petrasso, H.-S. Park, B. A. Remington, J. S. Ross, D. Ryu, D. Ryutov, T. G. White, B. Reville, F. Miniati, A. A. Schekochihin, D. Q. Lamb, D. H. Froula, and G. Gregori. "Laboratory Evidence of Dynamo Amplification of Magnetic Fields in a Turbulent Plasma," *Nature Communications* (February 2018), Nature Publishing Group.



# Towards a Definitive Model of Core-Collapse Supernova Explosions

PI Adam Burrows  
 INST Princeton University  
 AWARD INCITE  
 HOURS Theta: 1,750,000 Node-Hours

Core-collapse supernovae dramatically announce the death of massive stars and the birth of neutron stars. These violent explosions, which produce the highest densities of matter and energy in the universe, are responsible for creating most of the elements in nature. A fundamental theoretical understanding of such explosions is needed to advance research in nuclear and particle physics, and inform the interpretation of data from large-scale experiments. To shed light on this mysterious cosmological phenomenon, a research team led by Princeton University is using ALCF supercomputers to address whether and how 3D supernova explosion models differ from their 2D counterparts.

**CHALLENGE** Since the 1960s, there has been an agonizingly slow march towards demonstrating a robust mechanism of supernovae explosion. 2D simulations of supernovae have supported the theory that capturing a small fraction of the neutrinos emitted during collapse powers the explosions, but detailed 3D calculations proving this paradigm are lacking. With the power of leadership-class supercomputers and continued advances in software, researchers now have the capabilities to tackle this longstanding challenge in nuclear astrophysics.

**APPROACH** To carry out this study, the team is using FORNAX, their new highly scalable, 3D radiation-hydrodynamics code. By addressing the transport operator with an explicit method, the code significantly reduces the computational complexity and communication overhead of traditional multidimensional radiative transfer solutions by bypassing the need for global iterative solvers. The team is running FORNAX on Theta to carry out large-scale simulations aimed at determining if the neutrino mechanism is a robust



This visualization depicts the neutrino-driven roiling convection of the nuclear material behind the newly reenergized supernova shock wave that surrounds the newly birthed neutron star, just hundreds of milliseconds after shock revival. Image: Joseph A. Insley and Silvio Rizzi, Argonne National Laboratory; Adam Burrows and Hiroki Nagakura, Princeton University

explanation for supernova explosions and the stellar progenitor dependence of their observables.

**RESULTS** Building upon previous work, the team is using Theta to complete a suite of non-rotating 3D supernova models. The suite, which includes simulations of the collapse, bounce, and explosion of progenitor massive stars at several different solar masses, is the largest such 3D investigation in the history of core-collapse theory. The researchers are using these simulations for several studies, including an exploration of the impact of spatial resolution on the outcome and character of 3D supernova simulations. In another study, they are exploring the correlations between the time series and angular distributions of the neutrino and gravitational wave emissions and the real-time dynamics of the shock and the proto-neutron star core.

**IMPACT** The team's efforts to advance the fundamental theoretical understanding of supernova explosions will benefit ongoing research efforts to determine the origin of the elements in the universe, measure gravitational waves, and interpret laboratory nuclear reaction rate measurements in light of stellar nucleosynthesis.



# ALCF Projects

## 2019 INCITE

### BIOLOGICAL SCIENCES

#### Advancing Design and Structure Prediction of Proteins and Peptides

PI David Baker  
 INST University of Washington  
 HOURS Mira: 5,000,000 Node-Hours  
 Theta: 715,000 Node-Hours

#### Finite Difference Time Domain Simulations to Facilitate Early-Stage Human Cancer Detection

PI Allen Taflove  
 INST Northwestern University  
 HOURS Mira: 5,000,000 Node-Hours

### CHEMISTRY

#### Dynamic Nanocluster Catalysis

PI Anastassia Alexandrova  
 INST University of California, Los Angeles  
 HOURS Mira: 20,000,000 Node-Hours

### EARTH SCIENCE

#### Energy Exascale Earth System Model

PI Mark Taylor  
 INST Sandia National Laboratories  
 HOURS Theta: 2,354,000 Node-Hours  
 Summit: 700,000 Node-Hours

#### Extreme-Scale Simulations for Advanced Seismic Ground Motion and Hazard Modeling

PI Christine Goulet  
 INST University of Southern California  
 HOURS Mira: 1,620,000 Node-Hours  
 Summit: 134,000 Node-Hours

#### High-Resolution Climate Sensitivity and Prediction Simulations with the CESM

PI Gerald Meehl  
 INST NCAR  
 HOURS Mira: 13,250,000 Node-Hours

### ENERGY TECHNOLOGIES

#### Bringing NE Regulatory Workflow on a Path to Exascale: LES Validation for PANDA

PI Aleksandr Obabko  
 INST Argonne National Laboratory  
 HOURS Mira: 17,700,000 Node-Hours

### ENGINEERING

#### Crystal Plasticity from First Principles

PI Vasily Bulatov  
 INST Lawrence Livermore National Laboratory  
 HOURS Mira: 17,000,000 Node-Hours

#### DNS Reference Data for Turbulence Model Development on the Bachalo-Johnson ATB

PI Koen Hillewaert  
 INST Cenaero  
 HOURS Mira: 7,875,000 Node-Hours

#### Laminar-Turbulent Transition in Swept-Wing Boundary Layers

PI Lian Duan  
 INST Missouri University of Science and Technology  
 HOURS Mira: 6,000,000 Node-Hours

#### Large-Eddy Simulation of a Commercial Transport Aircraft Model

PI Parviz Moin  
 INST Stanford University  
 HOURS Mira: 24,000,000 Node-Hours

#### Shock-Induced Multi-Material Mixing

PI Sanjiva Lele  
 INST Stanford University  
 HOURS Mira: 4,500,000 Node-Hours

#### Towards Ultimate Rayleigh-Benard Convection

PI Janet Scheel  
 INST Occidental College  
 HOURS Mira: 12,500,000 Node-Hours

### MATERIALS SCIENCE

#### Advancing Electronic Stopping Simulation: From Solids to DNA

PI Yosuke Kanai  
 INST University of North Carolina at Chapel Hill  
 HOURS Mira: 4,180,000 Node-Hours  
 Theta: 1,000,000 Node-Hours

#### Materials and Interfaces for Organic and Hybrid Photovoltaics

PI Noa Marom  
 INST Carnegie Mellon University  
 HOURS Mira: 16,000,000 Node-Hours  
 Theta: 1,000,000 Node-Hours

#### Petascale Simulations for Layered Materials Genome

PI Aiichiro Nakano  
 INST University of Southern California  
 HOURS Mira: 12,500,000 Node-Hours

#### Predictive Simulations of Functional Materials

PI Paul Kent  
 INST Oak Ridge National Laboratory  
 HOURS Theta: 1,750,000 Node-Hours  
 Titan: 1,000,000 Node-Hours  
 Summit: 100,000 Node-Hours

#### Reactive Mesoscale Simulations of Tribological Interfaces

PI Subramanian Sankaranarayanan  
 INST Argonne National Laboratory  
 HOURS Theta: 650,000 Node-Hours



## PHYSICS

### Ab-initio Nuclear Structure and Nuclear Reactions

PI Gaute Hagen  
INST Oak Ridge National Laboratory  
HOURS Mira: 3,500,000 Node-Hours  
Theta: 650,000 Node-Hours  
Titan: 3,500,000 Node-Hours  
Summit: 364,000 Node-Hours

### Astrophysical Particle Accelerators: Magnetic Reconnection and Turbulence

PI Dmitri Uzdensky  
INST University of Colorado Boulder  
HOURS Mira: 6,750,000 Node-Hours

### Extreme-Scale Simulation of Supernovae and Magnetars from Realistic Progenitors

PI Sean Couch  
INST Michigan State University  
HOURS Mira: 9,375,000 Node-Hours  
Theta: 281,000 Node-Hours

### High-Fidelity Gyrokinetic Simulation of Tokamak and ITER Edge Physics

PI Choong-Seock Chang  
INST Princeton Plasma Physics Laboratory  
HOURS Theta: 1,500,000 Node-Hours  
Titan: 3,500,000 Node-Hours  
Summit: 1,050,000 Node-Hours

### Kinetic Simulation of FRC Stability and Transport

PI Sean Dettrick  
INST TAE Technologies, Inc.  
HOURS Theta: 750,000 Node-Hours

### Lattice QCD

PI Paul Mackenzie  
INST Fermilab  
HOURS Mira: 24,000,000 Node-Hours  
Theta: 1,600,000 Node-Hours

### N-Jettiness Subtraction for Precision Collider Phenomenology

PI Radja Boughezal  
INST Argonne National Laboratory  
HOURS Mira: 1,250,000 Node-Hours  
Theta: 400,000 Node-Hours

### Petascale Simulations of Kinetic Effects in IFE Plasmas

PI Frank Tsung  
INST University of California, Los Angeles  
HOURS Theta: 1,500,000 Node-Hours

### PLASM-IN-SILICO: HPC Modeling of High-Intensity Laser-Solid Interaction

PI Jean-Luc Vay  
INST Lawrence Berkeley National Laboratory  
HOURS Mira: 8,000,000 Node-Hours  
Theta: 600,000 Node-Hours

### Radiation Hydrodynamic Simulations of Massive Stars with Rotation

PI Lars Bildsten  
INST University of California, Santa Barbara  
HOURS Theta: 1,300,000 Node-Hours

### Towards a Definitive Model of Core-Collapse Supernova Explosions

PI Adam Burrows  
INST Princeton University  
HOURS Theta: 1,750,000 Node-Hours

## ALCC 2018-2019

*The 2018-2019 ALCC award allocations are shown in core-hours. In 2019, the allocation programs transitioned to awarding DOE computing awards in node-hours.*

## CHEMISTRY

### High-Fidelity Simulations of Flow and Heat Transfer During Motored Operation of an Internal Combustion Engine

PI Paul Fischer  
INST Argonne National Laboratory  
HOURS 30,000,000 Core-Hours

## COMPUTER SCIENCE

### Demonstration of the Scalability of Programming Environments by Simulating Multiscale Applications

PI Robert Voigt  
INST Leidos, Inc.  
HOURS 198,500,000 Core-Hours  
ALCF: 100M, OLCF: 78.5M, NERSC: 20M

### Portable Application Development for Next-Generation Supercomputer Architectures

PI T.P. Straatsma  
INST Oak Ridge National Laboratory  
HOURS 60,000,000 Core-Hours  
ALCF: 30M, OLCF: 30M

## EARTH SCIENCE

### Investigating the Impact of Improved Southern Ocean Processes in Antarctic-Focused Global Climate Simulations

PI Mark Petersen  
INST Los Alamos National Laboratory  
HOURS 105,000,000 Core-Hours  
ALCF: 35M, OLCF: 5M, NERSC: 65M

## ENERGY TECHNOLOGIES

### High-Fidelity Numerical Simulation of Wire-Wrapped Fuel Assemblies: Year 2

PI Aleksandr Obabko  
INST Argonne National Laboratory  
HOURS 83,500,000 Core-Hours

### High-Fidelity Simulation for Molten Salt Reactors: Enabling Innovation Through Petascale Computing

PI Elia Merzari  
INST Argonne National Laboratory  
HOURS 140,000,000 Core-Hours

### HPC4EnergyInnovation ALCC End-Station

PI Peter Nugent  
INST Lawrence Berkeley National Laboratory  
HOURS 170,000,000 Core-Hours  
ALCF: 20M, OLCF: 100M, NERSC: 50M

### Multiphase Flow Simulations of Nuclear Reactor Flows

PI Igor Bolotnov  
INST North Carolina State University  
HOURS 130,000,000 Core-Hours

## ENGINEERING

### Analysis and Mitigation of Dynamic Stall in Energy Machines

PI Anupam Sharma  
INST Iowa State University  
HOURS 51,500,000 Core-Hours

## MATERIALS SCIENCE

### Imaging and Controlling Elemental Contrast of Nanocluster in Intense X-ray Pulses

PI Phay Ho  
INST Argonne National Laboratory  
HOURS 90,000,000 Core-Hours

### Impact of Grain Boundary Defects on Hybrid Perovskite Solar Absorbers

PI Wissam Saidi  
INST University of Pittsburgh  
HOURS 20,000,000 Core-Hours

### Large-Scale Simulations of Heterogeneous Materials for Energy Conversion Applications

PI Giulia Galli  
INST The University of Chicago and Argonne National Laboratory  
HOURS 100,000,000 Core-Hours

### Modeling Fusion Plasma Facing Components

PI Brian Wirth  
INST Oak Ridge National Laboratory and University of Tennessee  
HOURS 165,000,000 Core-Hours  
ALCF: 80M, OLCF: 60M, NERSC: 25M

### Predictive Modeling and Machine Learning for Functional Nanoporous Materials

PI J. Ilja Siepmann  
INST University of Minnesota  
HOURS 58,000,000 Core-Hours  
ALCF: 42M, NERSC: 16M



## PHYSICS

**Emulating the Universe**

PI Katrin Heitmann  
 INST Argonne National Laboratory  
 HOURS 50,000,000 Core-Hours  
*ALCF: 10M, OLCF: 40M*

**Hadronic Light-by-Light Scattering and Vacuum Polarization Contributions to the Muon Anomalous Magnetic Moment from Lattice QCD with Chiral Fermions**

PI Thomas Blum  
 INST University of Connecticut  
 HOURS 162,000,000 Core-Hours

**Nucleon Structure and Electric Dipole Moments with Physical Chiral-Symmetric Quarks**

PI Sergey Syritsyn  
 INST RIKEN BNL Research Center  
 HOURS 50,000,000 Core-Hours

**Scaling LHC Proton-Proton Collision Simulations in the ATLAS Detector**

PI Eric Lancon  
 INST Brookhaven National Laboratory  
 HOURS 160,000,000 Core-Hours  
*ALCF: 80M, OLCF: 80M*

**Semileptonic B- and D-meson Form Factors with High Precision**

PI Ruth Van de Water  
 INST Fermi National Accelerator Laboratory  
 HOURS 247,000,000 Core-Hours

**Simulations of Laser Experiments to Study MHD Turbulence and Non-Thermal Charged Particles**

PI Petros Tzeferacos  
 INST The University of Chicago  
 HOURS 22,000,000 Core-Hours

## ALCC 2019-2020

## ENERGY TECHNOLOGIES

**High-Fidelity Physics Simulations for DOE and Industry Fast Spectrum Nuclear Reactor Systems**

PI Emily Shemon  
 INST Argonne National Laboratory  
 HOURS Theta: 880,000 Node-Hours

**Nuclear Energy Industry Validation of Nek5000: ALAIN and HYMERES**

PI Aleksander Obabko  
 INST Argonne National Laboratory  
 HOURS Theta: 340,000 Node-Hours

## ENGINEERING

**Towards Exascale Internal Combustion Engine Simulation with In-Situ Analysis**

PI Muhsin Ameen  
 INST Argonne National Laboratory  
 HOURS Theta: 630,000 Node-Hours

## MATERIALS SCIENCE

**Accelerated Catalyst Discovery from First Principles Simulations and Machine Learning**

PI Rajeev Surendran Assary  
 INST Argonne National Laboratory  
 HOURS Theta: 240,000 Node-Hours

**Predictive Modeling and Machine Learning for Functionally Nanoporous Materials**

PI J. Ilja Siepmann  
 INST University of Minnesota  
 HOURS Theta: 620,000 Node-Hours  
 Cori: 200,000 Node-Hours

**Supercomputing for Automotive High-Temperature Alloy Design**

PI Dongwon Shin  
 INST Oak Ridge National Laboratory  
 HOURS Theta: 350,000 Node-Hours  
 Summit: 391,000 Node-Hours

## PHYSICS

**Low Energy Neutrino-Nucleus Interactions**

PI Saori Pastore  
 INST Washington University in St. Louis  
 HOURS Theta: 390,000 Node-Hours

**Modeling the Response of Fusion Plasma Components**

PI Brian Wirth  
 INST Oak Ridge National Laboratory  
 HOURS Theta: 200,000 Node-Hours  
 Summit: 250,000 Node-Hours

**Neutrino Flux, Energy Deposition and Radiological Studies for the DUNE-LBNF Beamline**

PI Igor Rakhno  
 INST Fermi National Accelerator Laboratory  
 HOURS Theta: 450,000 Node-Hours

**The Next Leap Forward in LSST Sky Simulations**

PI Katrin Heitmann  
 INST Argonne National Laboratory  
 HOURS Theta: 400,000 Node-Hours

**Particle Acceleration in Plasma Jets: From Astrophysics to the Laboratory**

PI Paulo E. Alves  
 INST SLAC National Accelerator Laboratory  
 HOURS Theta: 1,000,000 Node-Hours

**Semileptonic B- and D-Meson Form Factors with High Precision**

PI Aida El-Khadra  
 INST University of Illinois at Urbana-Champaign  
 HOURS Theta: 400,000 Node-Hours  
 Cori: 330,000 Node-Hours

## ALCF Data Science Program

**Constructing and Navigating Polymorphic Landscapes of Molecular Crystals**

PI Alexandre Tkatchenko  
 INST University of Luxembourg  
 HOURS Mira: 3,125,000 Node-Hours  
 Theta: 125,000 Node-Hours

**Data-Driven Materials Discovery for Optoelectronic Applications**

PI Jacqueline Cole  
 INST University of Cambridge  
 HOURS Theta: 468,750 Node-Hours

**Deep Learning at Scale for Multimessenger Astrophysics Through the NCSA-Argonne Collaboration**

PI Eliu Huerta  
 INST University of Illinois at Urbana-Champaign  
 HOURS Theta: 156,250 Node-Hours

**Developing High-Fidelity Dynamic and Ultrafast X-ray Imaging Tools for APS-Upgrade**

PI Jin Wang  
 INST Argonne National Laboratory  
 HOURS Theta: 125,000 Node-Hours

**Enabling Multiscale Physics for Industrial Design Using Deep Learning Networks**

PI Rathakrishnan Bhaskaran  
 INST GE Global Research  
 HOURS Theta: 93,750 Node-Hours

**Machine Learning Magnetic Properties of van der Waals Heterostructures**

PI Efthimios Kaxiras  
 INST Harvard University  
 HOURS Theta: 156,250 Node-Hours

**Realistic Simulations of the LSST Survey at Scale**

PI Katrin Heitmann  
 INST Argonne National Laboratory  
 HOURS Theta: 218,750 Node-Hours

**X-ray Microscopy of Extended 3D Objects: Scaling Towards the Future**

PI Chris Jacobsen  
 INST Argonne National Laboratory and Northwestern University  
 HOURS Theta: 187,500 Node-Hours

## Aurora Early Science Program

**Accelerated Deep Learning Discovery in Fusion Energy Science**

PI William Tang  
 INST Princeton Plasma Physics Laboratory

**Dark Sky Mining**

PI Salman Habib  
 INST Argonne National Laboratory



#### Data Analytics and Machine Learning for Exascale Computational Fluid Dynamics

PI Kenneth Jansen  
INST University of Colorado Boulder

#### Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience

PI Nicola Ferrier  
INST Argonne National Laboratory

#### Exascale Computational Catalysis

PI David Bross  
INST Argonne National Laboratory

#### Extending Moore's Law Computing with Quantum Monte Carlo

PI Anouar Benali  
INST Argonne National Laboratory

#### Extreme-Scale Cosmological Hydrodynamics

PI Katrin Heitmann  
INST Argonne National Laboratory

#### Extreme-Scale In-Situ Visualization and Analysis of Fluid-Structure-Interaction Simulations

PI Amanda Randles  
INST Duke University and  
Oak Ridge National Laboratory

#### Extreme-Scale Unstructured Adaptive CFD

PI Kenneth Jansen  
INST University of Colorado Boulder

#### High-Fidelity Simulation of Fusion Reactor Boundary Plasmas

PI C.S. Chang  
INST Princeton Plasma Physics Laboratory

#### Machine Learning for Lattice Quantum Chromodynamics

PI William Detmold  
INST Massachusetts Institute of Technology

#### Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials

PI Noa Marom  
INST Carnegie Mellon University

#### NWChemEx: Tackling Chemical, Materials, and Biochemical Challenges in the Exascale Era

PI Thom Dunning  
INST Pacific Northwest National Laboratory

#### Simulating and Learning in the ATLAS Detector at the Exascale

PI James Proudfoot  
INST Argonne National Laboratory

#### Virtual Drug Response Prediction

PI Rick Stevens  
INST Argonne National Laboratory

## 2019 Director's Discretionary

*The following list provides a sampling of the many Director's Discretionary projects at the ALCF.*

### BIOLOGICAL SCIENCES

#### Computational Analysis of Brain Connectomes for Alzheimer's Disease

PI Jiok Cha  
INST Columbia University  
HOURS Theta: 46,875 Node-Hours

#### Free Energy Landscapes of Membrane Proteins

PI Benoît Roux  
INST The University of Chicago  
HOURS Theta: 78,125 Node-Hours

### CHEMISTRY

#### Accelerated Catalyst Discovery from First Principles Simulations and Machine Learning

PI Rajeev Surendran Assary  
INST Argonne National Laboratory  
HOURS Theta: 135,175 Node-Hours

#### Ionic Liquid as a Potential Electrolyte of High-Performance Lithium Ion Battery

PI Zhengcheng Zhang  
INST Argonne National Laboratory  
HOURS Theta: 93,750 Node-Hours

### COMPUTER SCIENCE

#### MPICH

PI Kenneth Raffanetti  
INST Argonne National Laboratory  
HOURS Theta: 281,250 Node-Hours

#### Simulating Realistic Quantum Computers

PI Hal Finkel  
INST Argonne National Laboratory  
HOURS Theta: 250,000 Node-Hours

#### Workload Interference Analysis on Theta

PI Kenneth Raffanetti  
INST Illinois Institute of Technology  
HOURS Theta: 65,630 Node-Hours

### EARTH SCIENCE

#### Extreme-Scale Simulations for Advanced Seismic Ground Motion and Hazard Modeling

PI Christine Goulet  
INST University of Southern California  
HOURS Theta: 111,250 Node-Hours

#### Terrestrial Ecosystem Carbon Cycle of the Conterminous U.S.

PI Jinxun Liu  
INST United States Geological Survey  
HOURS Theta: 251,000 Node-Hours

### ENERGY TECHNOLOGIES

#### Thermal Hydraulic Simulations for the Versatile Test Reactor

PI Dillon Shaver  
INST Argonne National Laboratory  
HOURS Mira: 578,440 Node-Hours

#### Validation of Two-Phase Flow Models and Critical Heat Flux Prediction for the Highly-Scalable CFD Code Nek-2P

PI Surya Vegendla  
INST Argonne National Laboratory  
HOURS Mira: 437,500 Node-Hours

### ENGINEERING

#### Data Analysis of Turbulent Channel Flow at High Reynolds Number

PI Robert Moser  
INST University of Texas at Austin  
HOURS Mira: 937,500 Node-Hours

#### High-Fidelity Simulation of Supersonic Turbulent Flow-Structure Interaction and Mixing

PI Ivan Bermejo Moreno  
INST University of Southern California  
HOURS Theta: 207,570 Node-Hours

#### Investigation of a Low Octane Gasoline Fuel for a Heavy-Duty Diesel Engine in a Low-Temperature Combustion Regime

PI Sibendu Som  
INST Argonne National Laboratory  
HOURS Mira: 1,875,000 Node-Hours

#### MaLTESE (Machine Learning Tool for Engine Simulations and Experiments)

PI Shashi Aithal, Prasanna Balaprakash  
INST Argonne National Laboratory  
HOURS Theta: 2,000 Node-Hours

#### Multiphase Simulations of Nuclear Reactor Thermal Hydraulics

PI Igor A. Bolotnov  
INST North Carolina State University  
HOURS Mira: 625,000 Node-Hours

#### Parallel Hierarchic Adaptive Stabilized Transient Analysis

PI Kenneth Jansen  
INST University of Colorado Boulder  
HOURS Mira: 3,373,628 Node-Hours

#### Wall-Resolved Simulations of Canonical Wall Bounded Flows

PI Ramesh Balakrishnan  
INST Argonne National Laboratory  
HOURS Mira: 2,531,250 Node-Hours



## MATERIALS SCIENCE

**Large-Scale Real-Space Electronic Structure Calculations for Understanding Energetics of Complex Defects in Materials**

PI Phani Sudheer Motamarri  
 INST University of Michigan  
 HOURS Theta: 93,750 Node-Hours

**Optimizing Superconductor Transport Properties through Large-Scale Simulation**

PI Andreas Glatz  
 INST Argonne National Laboratory  
 HOURS Cooley: 12,000 Node-Hours

**Phase Transitions in Water-Ice-Vapor System**

PI Subramanian Sankaranarayanan  
 INST Argonne National Laboratory  
 HOURS Mira: 625,000 Node-Hours

**Rational Design of Ultrastrong Composites**

PI Hendrik Heinz  
 INST University of Colorado Boulder  
 HOURS Mira: 206,250 Node-Hours

**Structure and Properties of Grain Boundaries in Materials for Energy Applications**

PI Wissam Saidi  
 INST University of Pittsburgh  
 HOURS Mira: 750,000 Node-Hours

## PHYSICS

**Extreme-Scale Cosmology**

PI Katrin Heitmann  
 INST Argonne National Laboratory  
 HOURS Mira: 50,000,000 Node-Hours

**MARS Energy Deposition and Neutrino Flux Simulations**

PI Nikolai V. Mokhov  
 INST Fermilab  
 HOURS Theta: 65,630 Node-Hours

**Pearl Necklace**

PI Michael David Schneider  
 INST Imperial College London  
 HOURS Theta: 258,860 Node-Hours

**Pion and Kaon Quark-Gluon Structure from Lattice QCD**

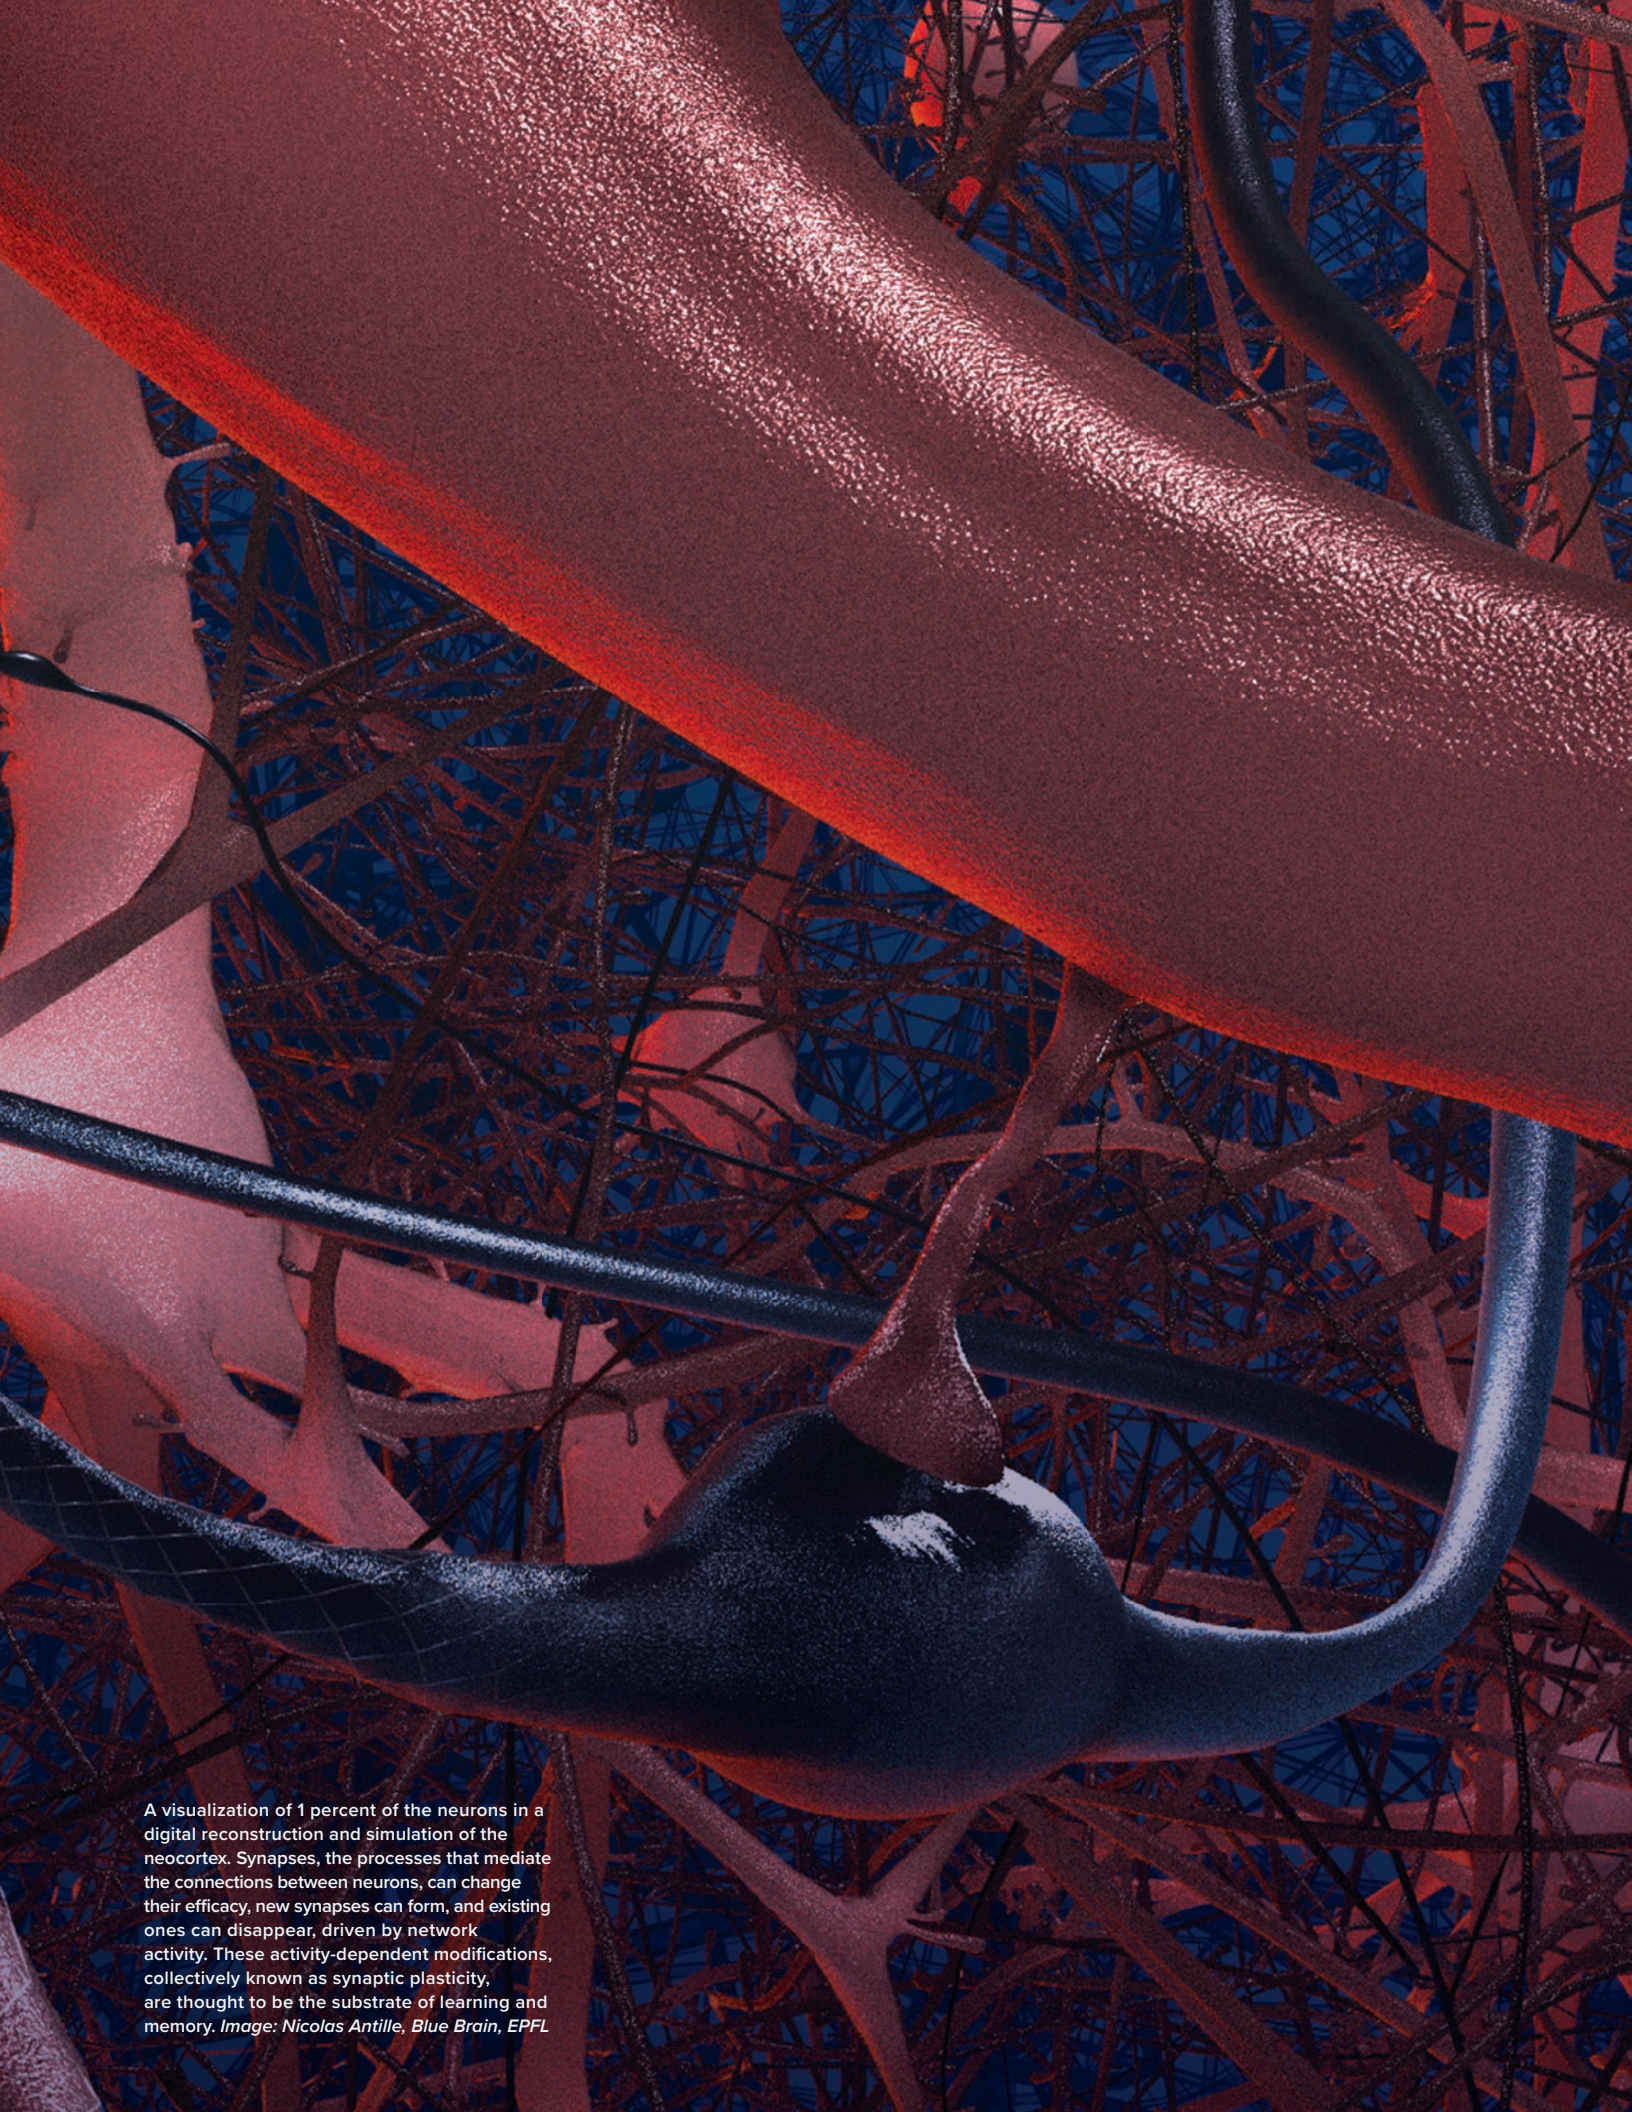
PI Ian Cloet  
 INST Argonne National Laboratory  
 HOURS Theta: 65,925 Node-Hours

**Unveiling the 3D Physics Behind Compact Ultrahigh Flux Neutron Sources**

PI Frederico Fiuza  
 INST SLAC National Accelerator Laboratory  
 HOURS Mira: 3,125,000 Node-Hours







A visualization of 1 percent of the neurons in a digital reconstruction and simulation of the neocortex. Synapses, the processes that mediate the connections between neurons, can change their efficacy, new synapses can form, and existing ones can disappear, driven by network activity. These activity-dependent modifications, collectively known as synaptic plasticity, are thought to be the substrate of learning and memory. Image: Nicolas Antille, Blue Brain, EPFL



## About the Argonne Leadership Computing Facility

---

Argonne's Leadership Computing Facility Division operates the Argonne Leadership Computing Facility (ALCF) as part of the U.S. Department of Energy's effort to provide leadership-class computing resources to the scientific community. The ALCF is supported by the DOE Office of Science, Advanced Scientific Computing Research (ASCR) program.

## About Argonne National Laboratory

---

Argonne is a U.S. Department of Energy Laboratory managed by UChicago Argonne, LLC, under contract DE-AC02-06CH11357. The Laboratory's main facility is outside of Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne and its pioneering science and technology programs, visit [www.anl.gov](http://www.anl.gov).

## Availability of this Report (ANL/ALCF-19/2)

---

Online Access: U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available for free via DOE's SciTech Connect (<http://www.osti.gov/scitech/>).

Reports not in digital format may be purchased by the public from the National Technical Information Service (NTIS):

U.S. Department of Commerce  
National Technical Information Service  
5301 Shawnee Rd.  
Alexandria, VA 22312

phone | 800.553.NTIS (6847) or 703.605.6000  
fax | 703.605.6900  
[orders@ntis.gov](mailto:orders@ntis.gov)  
[www.ntis.gov](http://www.ntis.gov)

Reports not in digital format are available to DOE and DOE contractors from the Office of Scientific and Technical Information (OSTI):

U.S. Department of Energy  
Office of Scientific and Technical Information  
P.O. Box 62  
Oak Ridge, TN 37831-0062

phone | 865.576.8401  
fax | 865.576.5728  
[reports@adonis.osti.gov](mailto:reports@adonis.osti.gov)  
[www.osti.gov](http://www.osti.gov)

ALCF Leadership: Michael E. Papka (Division Director), Jini Ramprakash (Deputy Division Director), Mark Fahey (Director of Operations), Katherine Riley (Director of Science), and Kalyan Kumaran (Director of Technology)

Editorial Team: Beth Cerny, Jim Collins, Nils Heinonen, Hayley Kim, and Laura Wolf

Design and production: Sandbox Studio, Chicago

Disclaimer: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor UChicago Argonne, LLC, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, Argonne National Laboratory, or UChicago Argonne, LLC.









## CONTACT

[media@alcf.anl.gov](mailto:media@alcf.anl.gov)

[alcf.anl.gov](http://alcf.anl.gov)



Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.

